

US007293237B1

(12) **United States Patent**  
**Knight et al.**

(10) **Patent No.:** **US 7,293,237 B1**  
(45) **Date of Patent:** **Nov. 6, 2007**

(54) **GRAPHICAL USER INTERFACE FOR CREATING A CUSTOMIZED STORAGE LAYOUT FOR AN APPLICATION**

6,542,972 B2 \* 4/2003 Ignatius et al. .... 711/154  
6,629,158 B1 \* 9/2003 Brant et al. .... 710/10

(75) Inventors: **Margaret E. Knight**, Seattle, WA (US);  
**Ko Wang**, Cupertino, CA (US);  
**Yateendra Kulkarni**, Sunnyvale, CA (US)

(Continued)

**OTHER PUBLICATIONS**

(73) Assignee: **Symantec Operating Corporation**,  
Cupertino, CA (US)

U.S. Appl. No. 10/388,193, filed Mar. 2003, Knight.

(Continued)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 826 days.

*Primary Examiner*—Tadesse Hailu  
(74) *Attorney, Agent, or Firm*—Meyertons Hood Kivlin Kowert & Goetzl, P.C.; Jeffrey C. Hood

(21) Appl. No.: **10/641,128**

(57) **ABSTRACT**

(22) Filed: **Aug. 14, 2003**

**Related U.S. Application Data**

(63) Continuation-in-part of application No. 10/327,561, filed on Dec. 20, 2002, now Pat. No. 7,043,619.

(60) Provisional application No. 60/348,870, filed on Jan. 14, 2002.

(51) **Int. Cl.**  
**G06F 3/00** (2006.01)

(52) **U.S. Cl.** ..... **715/735**; 715/771; 715/764;  
711/170; 711/171

(58) **Field of Classification Search** ..... 715/762–764,  
715/771, 700, 735, 736, 853; 711/170, 171;  
710/8, 15; 709/223, 224

See application file for complete search history.

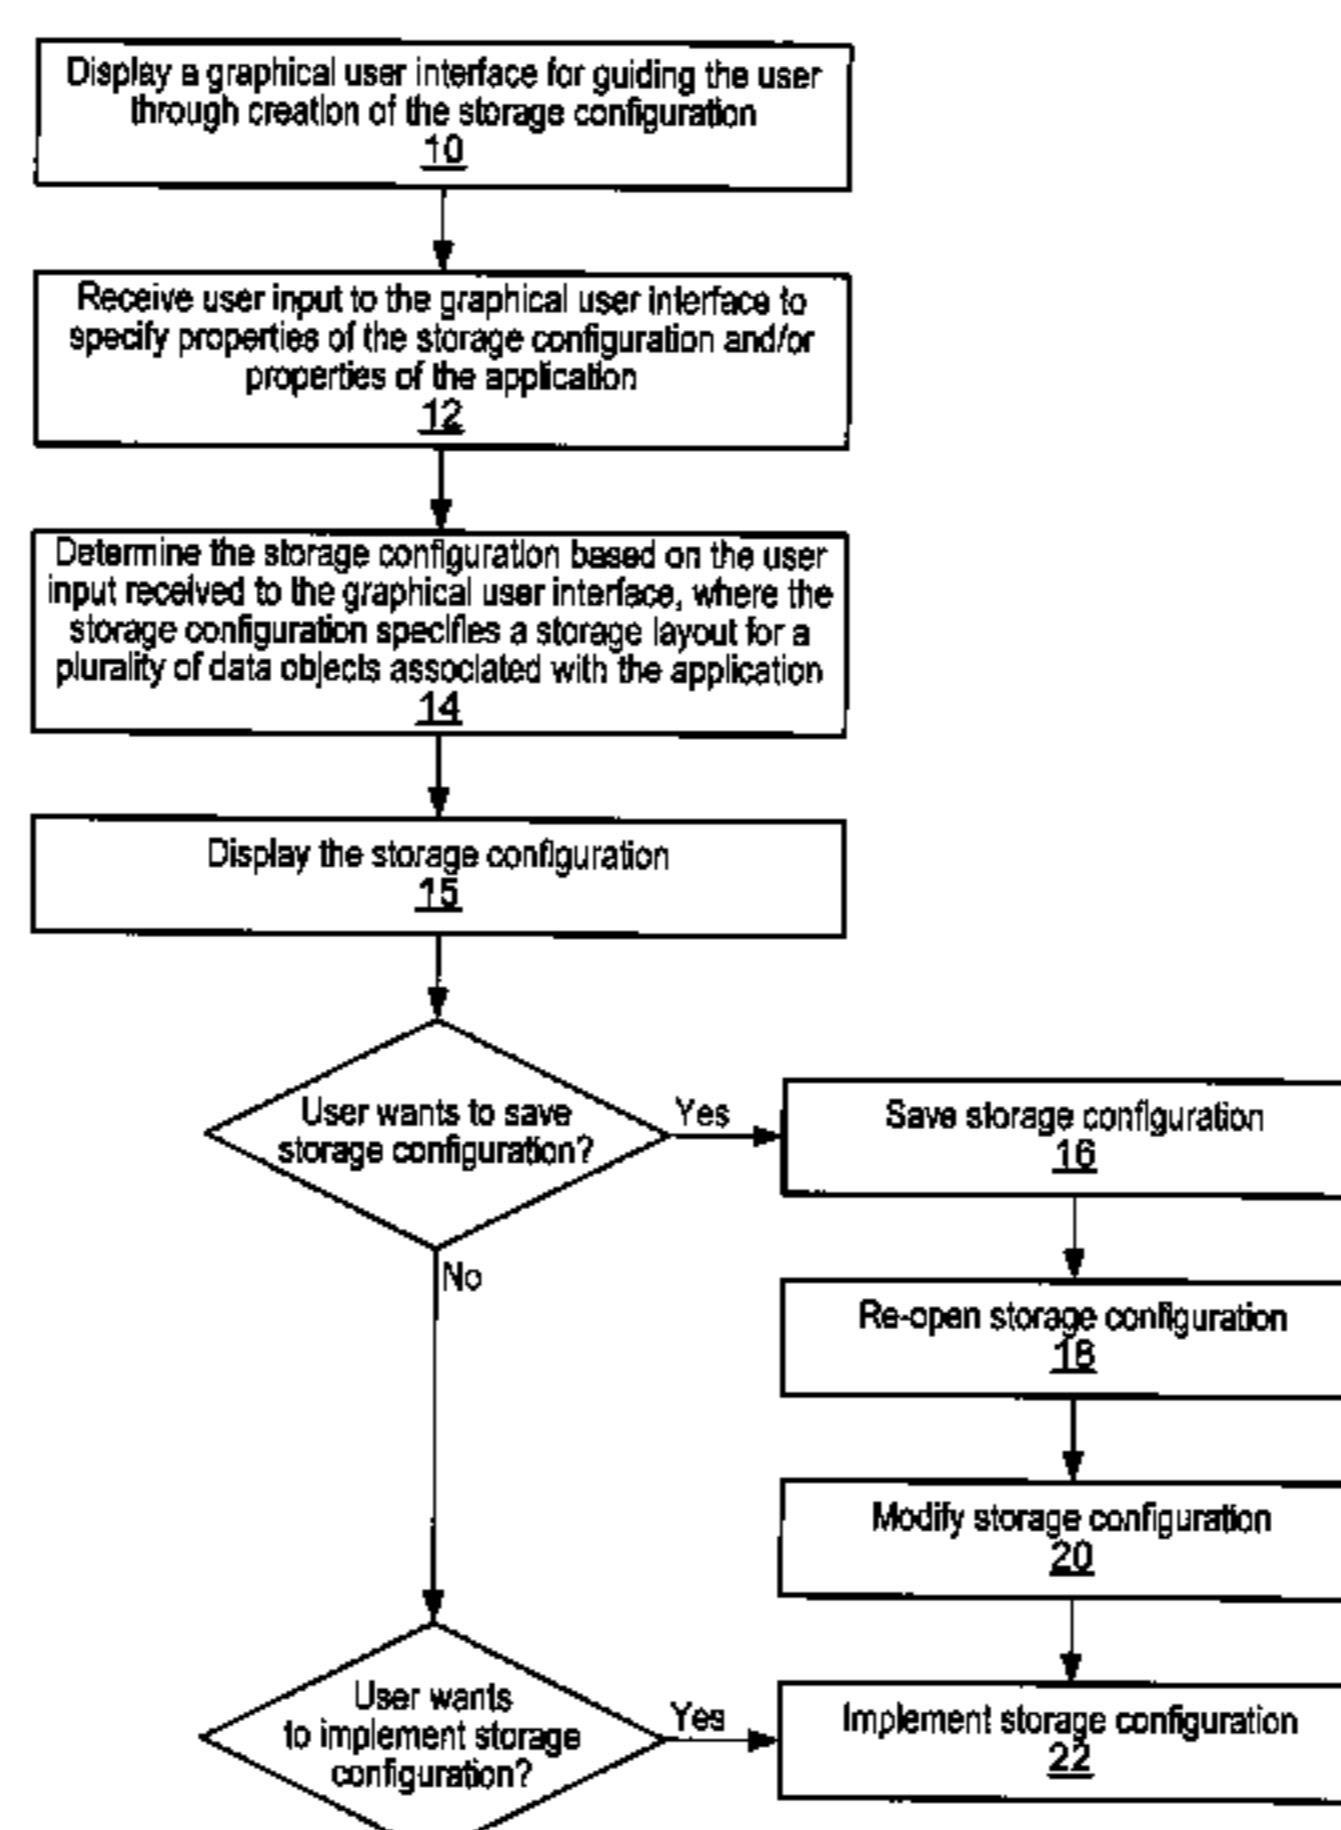
(56) **References Cited**

**U.S. PATENT DOCUMENTS**

5,890,204 A \* 3/1999 Ofer et al. .... 711/111  
5,951,687 A \* 9/1999 Chan et al. .... 713/2  
6,347,366 B1 2/2002 Cousins  
6,377,939 B1 \* 4/2002 Young ..... 705/34  
6,519,679 B2 2/2003 Devireddy et al.  
6,538,669 B1 3/2003 Lagueux, Jr. et al.

A system and method for determining a storage configuration for an application, where the storage configuration specifies a storage layout for data objects associated with the application. A Storage Configurator program may display a graphical user interface (GUI) for guiding the user through creation of the storage configuration. User input to specify properties of the storage configuration and/or properties of the application may be received to the graphical user interface. The Storage Configurator program may determine a storage configuration for the application, based on the user input received to the graphical user interface. In one embodiment, the Storage Configurator may provide an option allowing the user to save the storage configuration, e.g., in a file or database. Saving the storage configuration may allow the user to utilize the storage configuration at a later time. In one embodiment, the Storage Configurator may also provide an option allowing the user to implement the storage configuration. Implementing the storage configuration may comprise storing the plurality of data objects associated with the application on one or more storage devices according to the storage layout specified by the storage configuration. For example, the user may simply click a button to cause the data objects to be stored according to the storage layout.

**30 Claims, 29 Drawing Sheets**



# US 7,293,237 B1

Page 2

---

## U.S. PATENT DOCUMENTS

6,654,797 B1 11/2003 Kamper  
6,751,739 B1\* 6/2004 Verdun ..... 711/170  
2002/0128815 A1\* 9/2002 Merchant et al. .... 704/2  
2002/0129216 A1 9/2002 Collins  
2002/0133669 A1\* 9/2002 Devireddy et al. .... 711/114  
2002/0198888 A1\* 12/2002 Young ..... 707/100  
2003/0033398 A1\* 2/2003 Carlson et al. .... 709/223  
2003/0046270 A1 3/2003 Leung et al.

2003/0074235 A1\* 4/2003 Gregory ..... 705/4  
2003/0076349 A1\* 4/2003 Slaby ..... 345/735  
2006/0006905 A1\* 1/2006 Chou et al. .... 326/40

## OTHER PUBLICATIONS

“Performance and Capacity Planning Solutions”, © 1999 Compaq  
Computer Cooperation.

\* cited by examiner

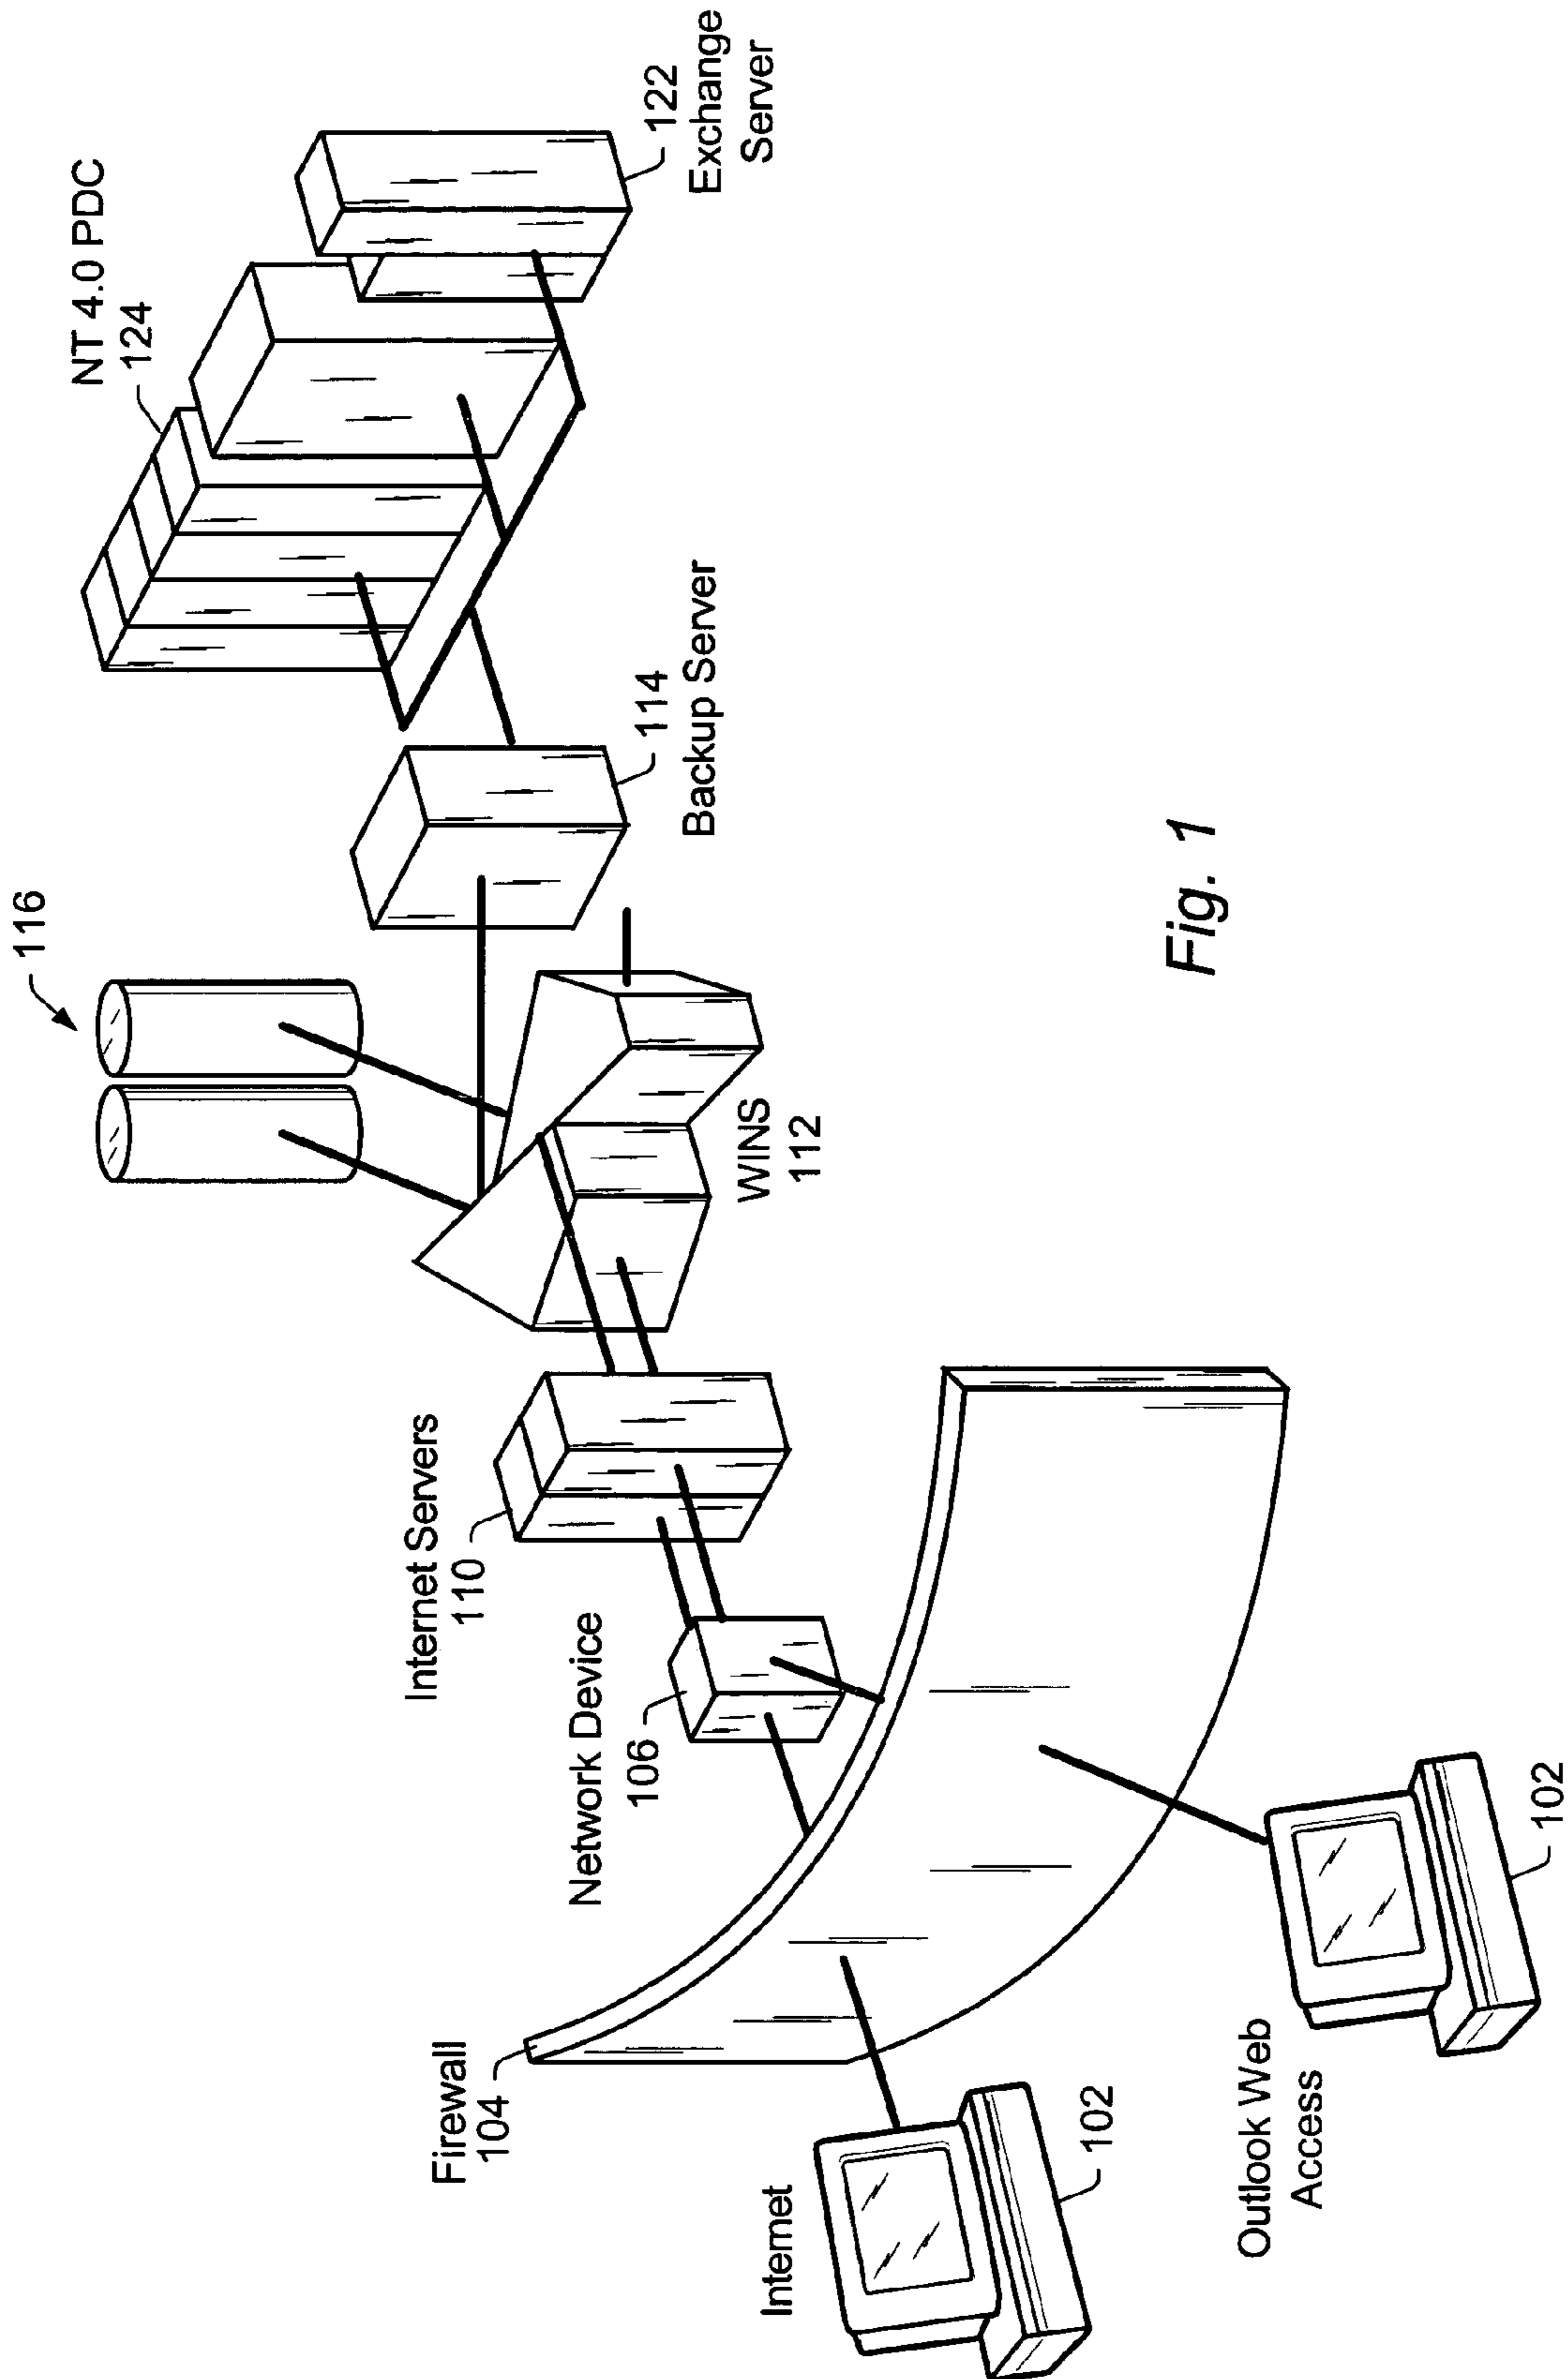


Fig. 1

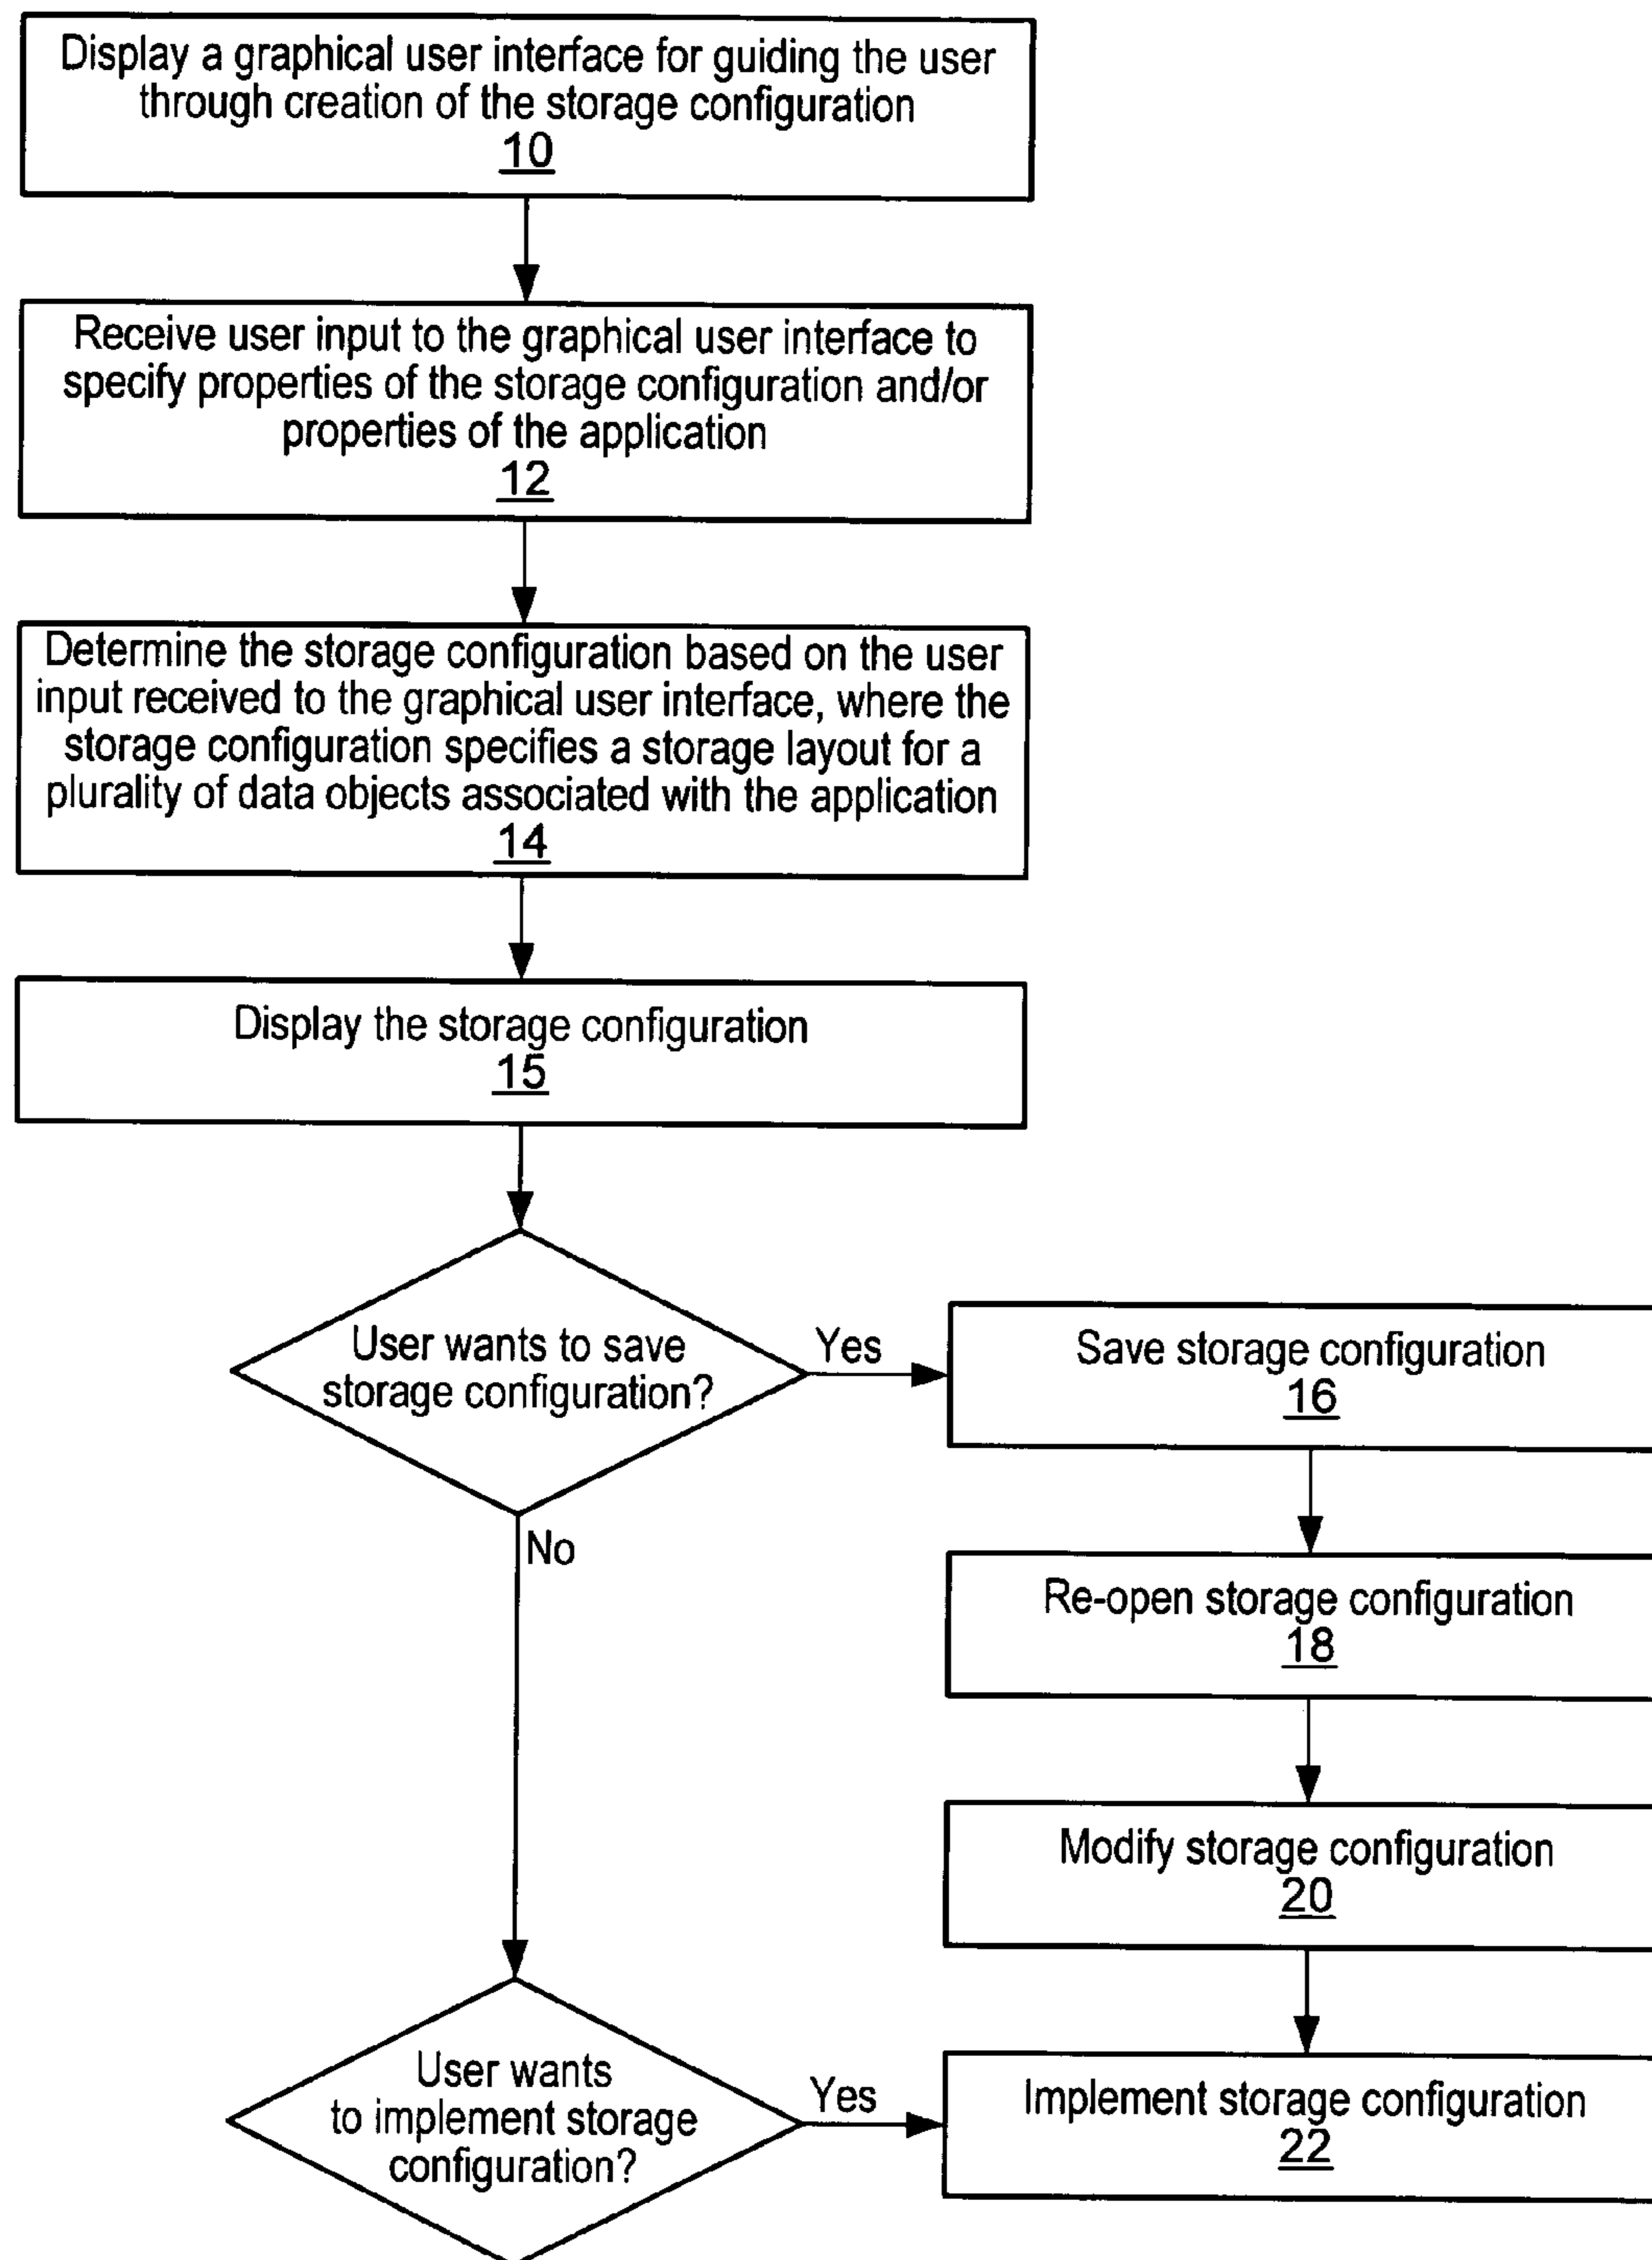


Fig. 2

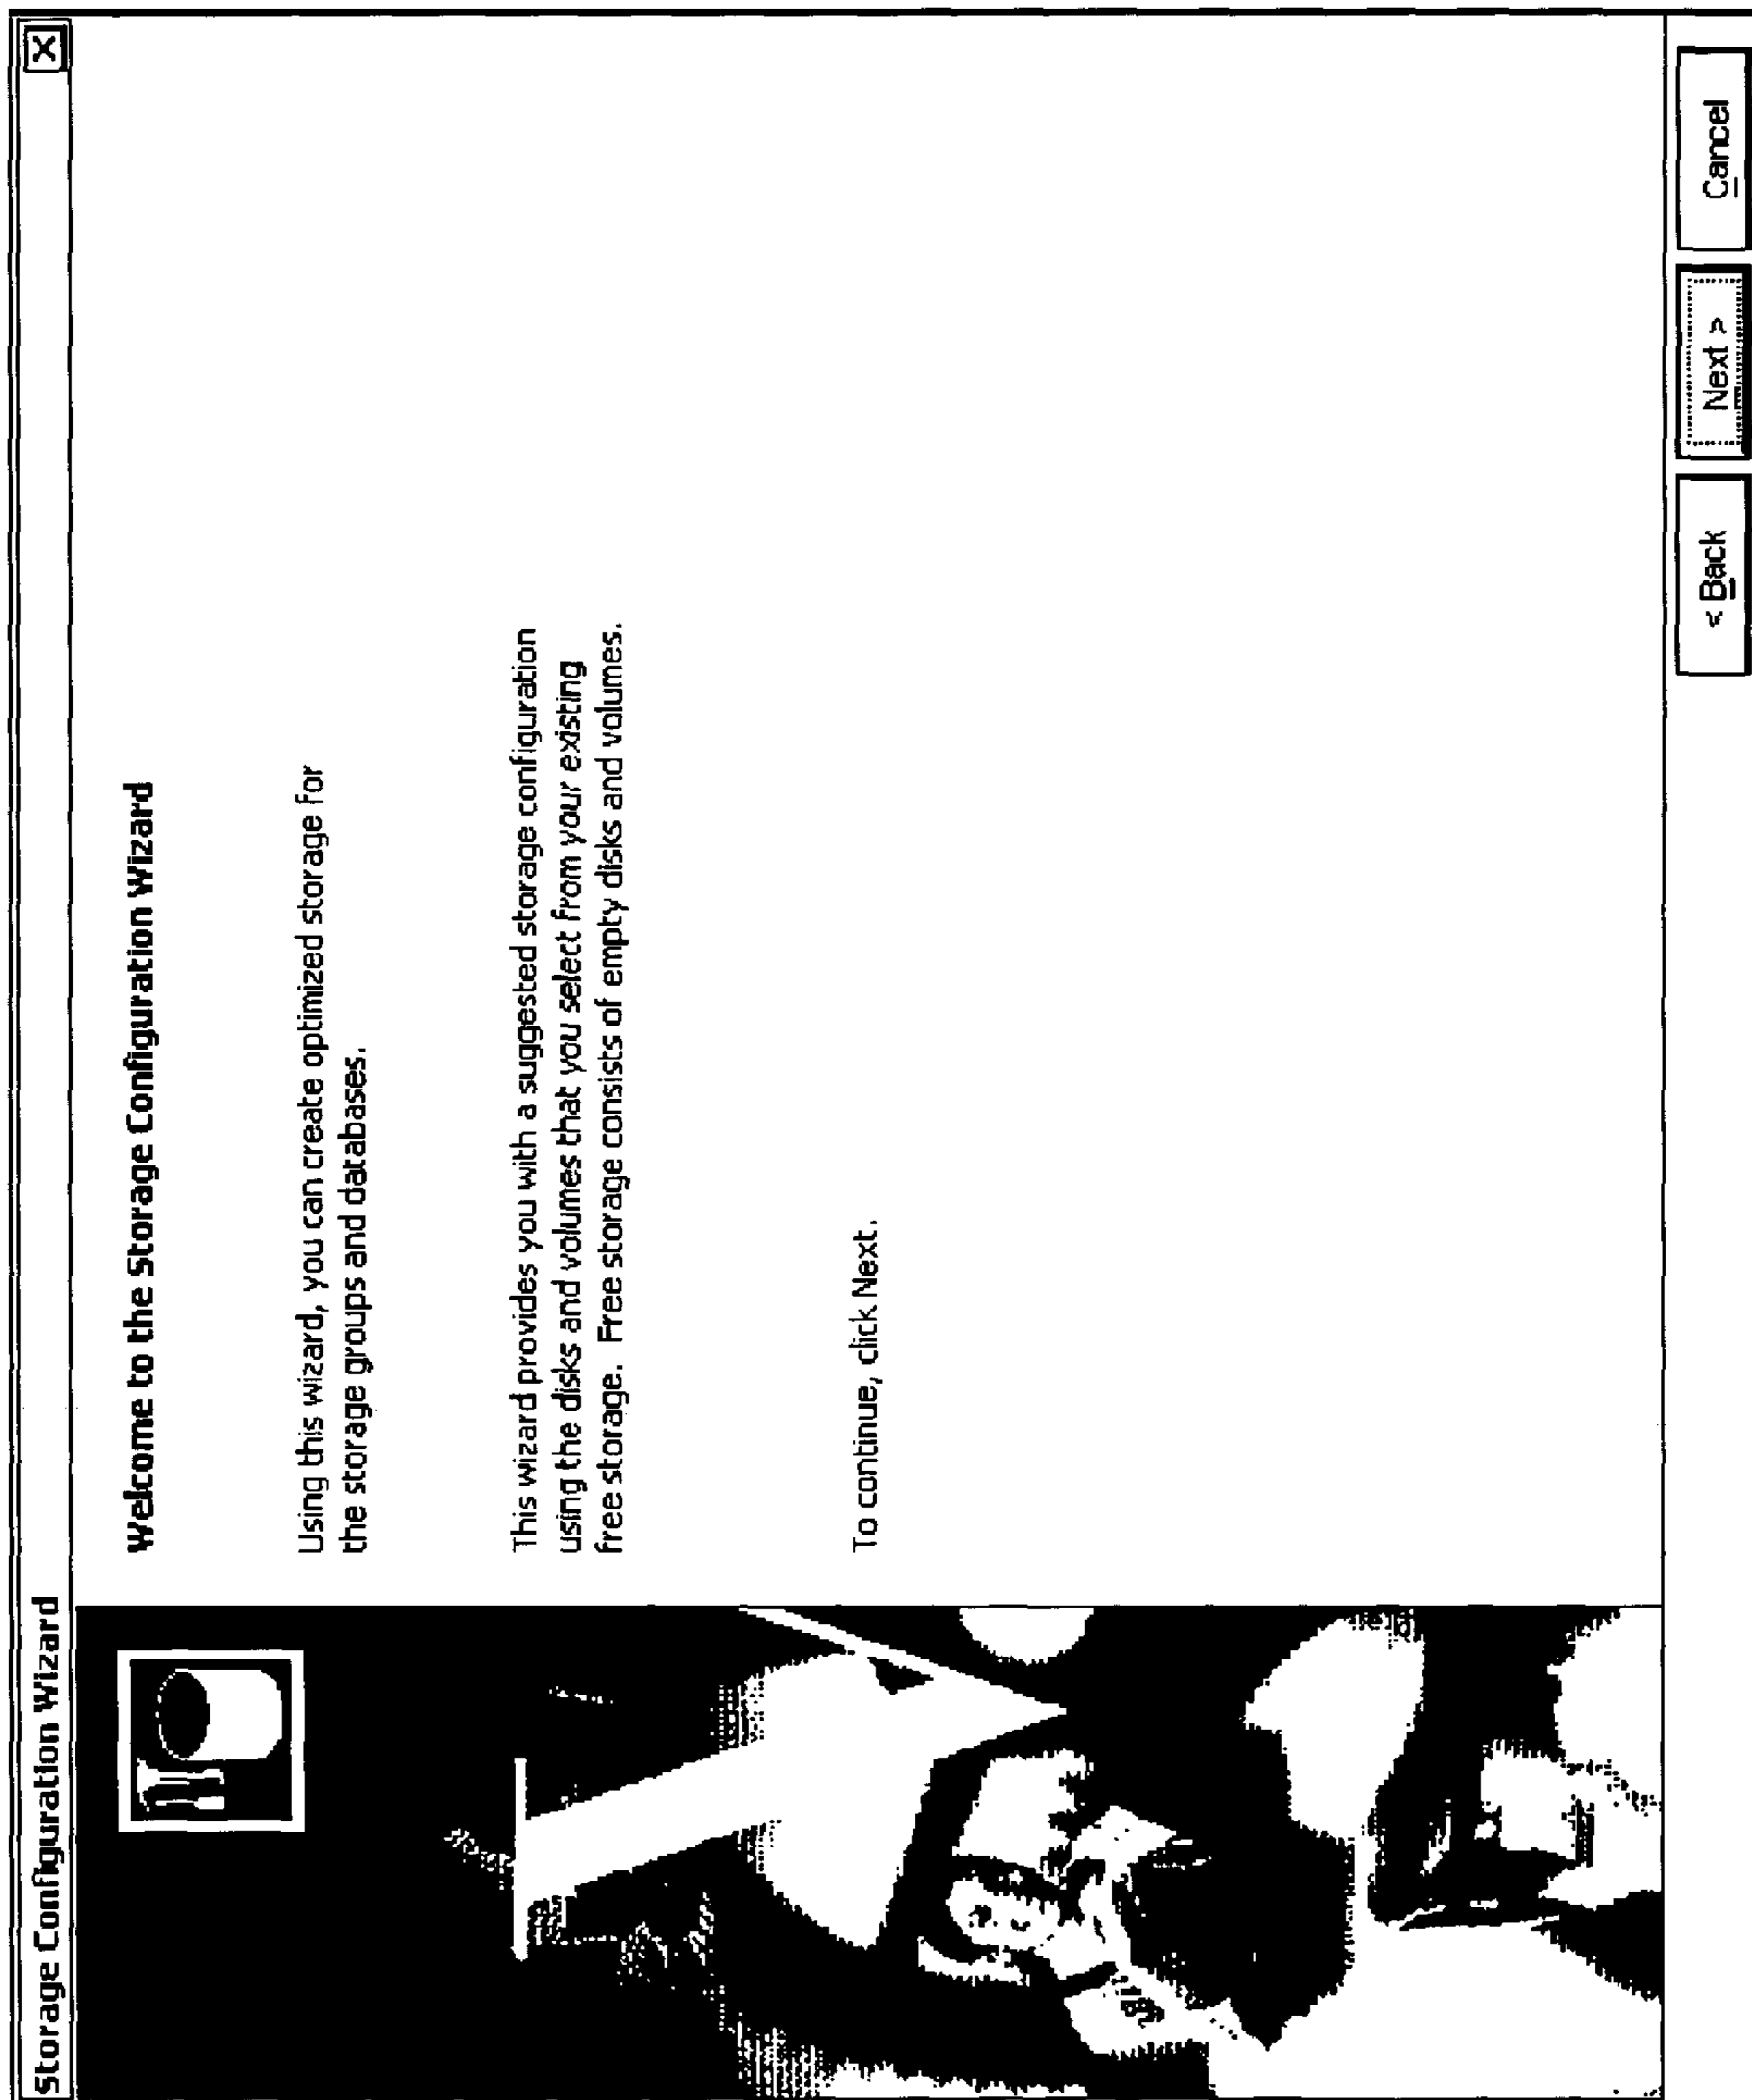


Fig. 3



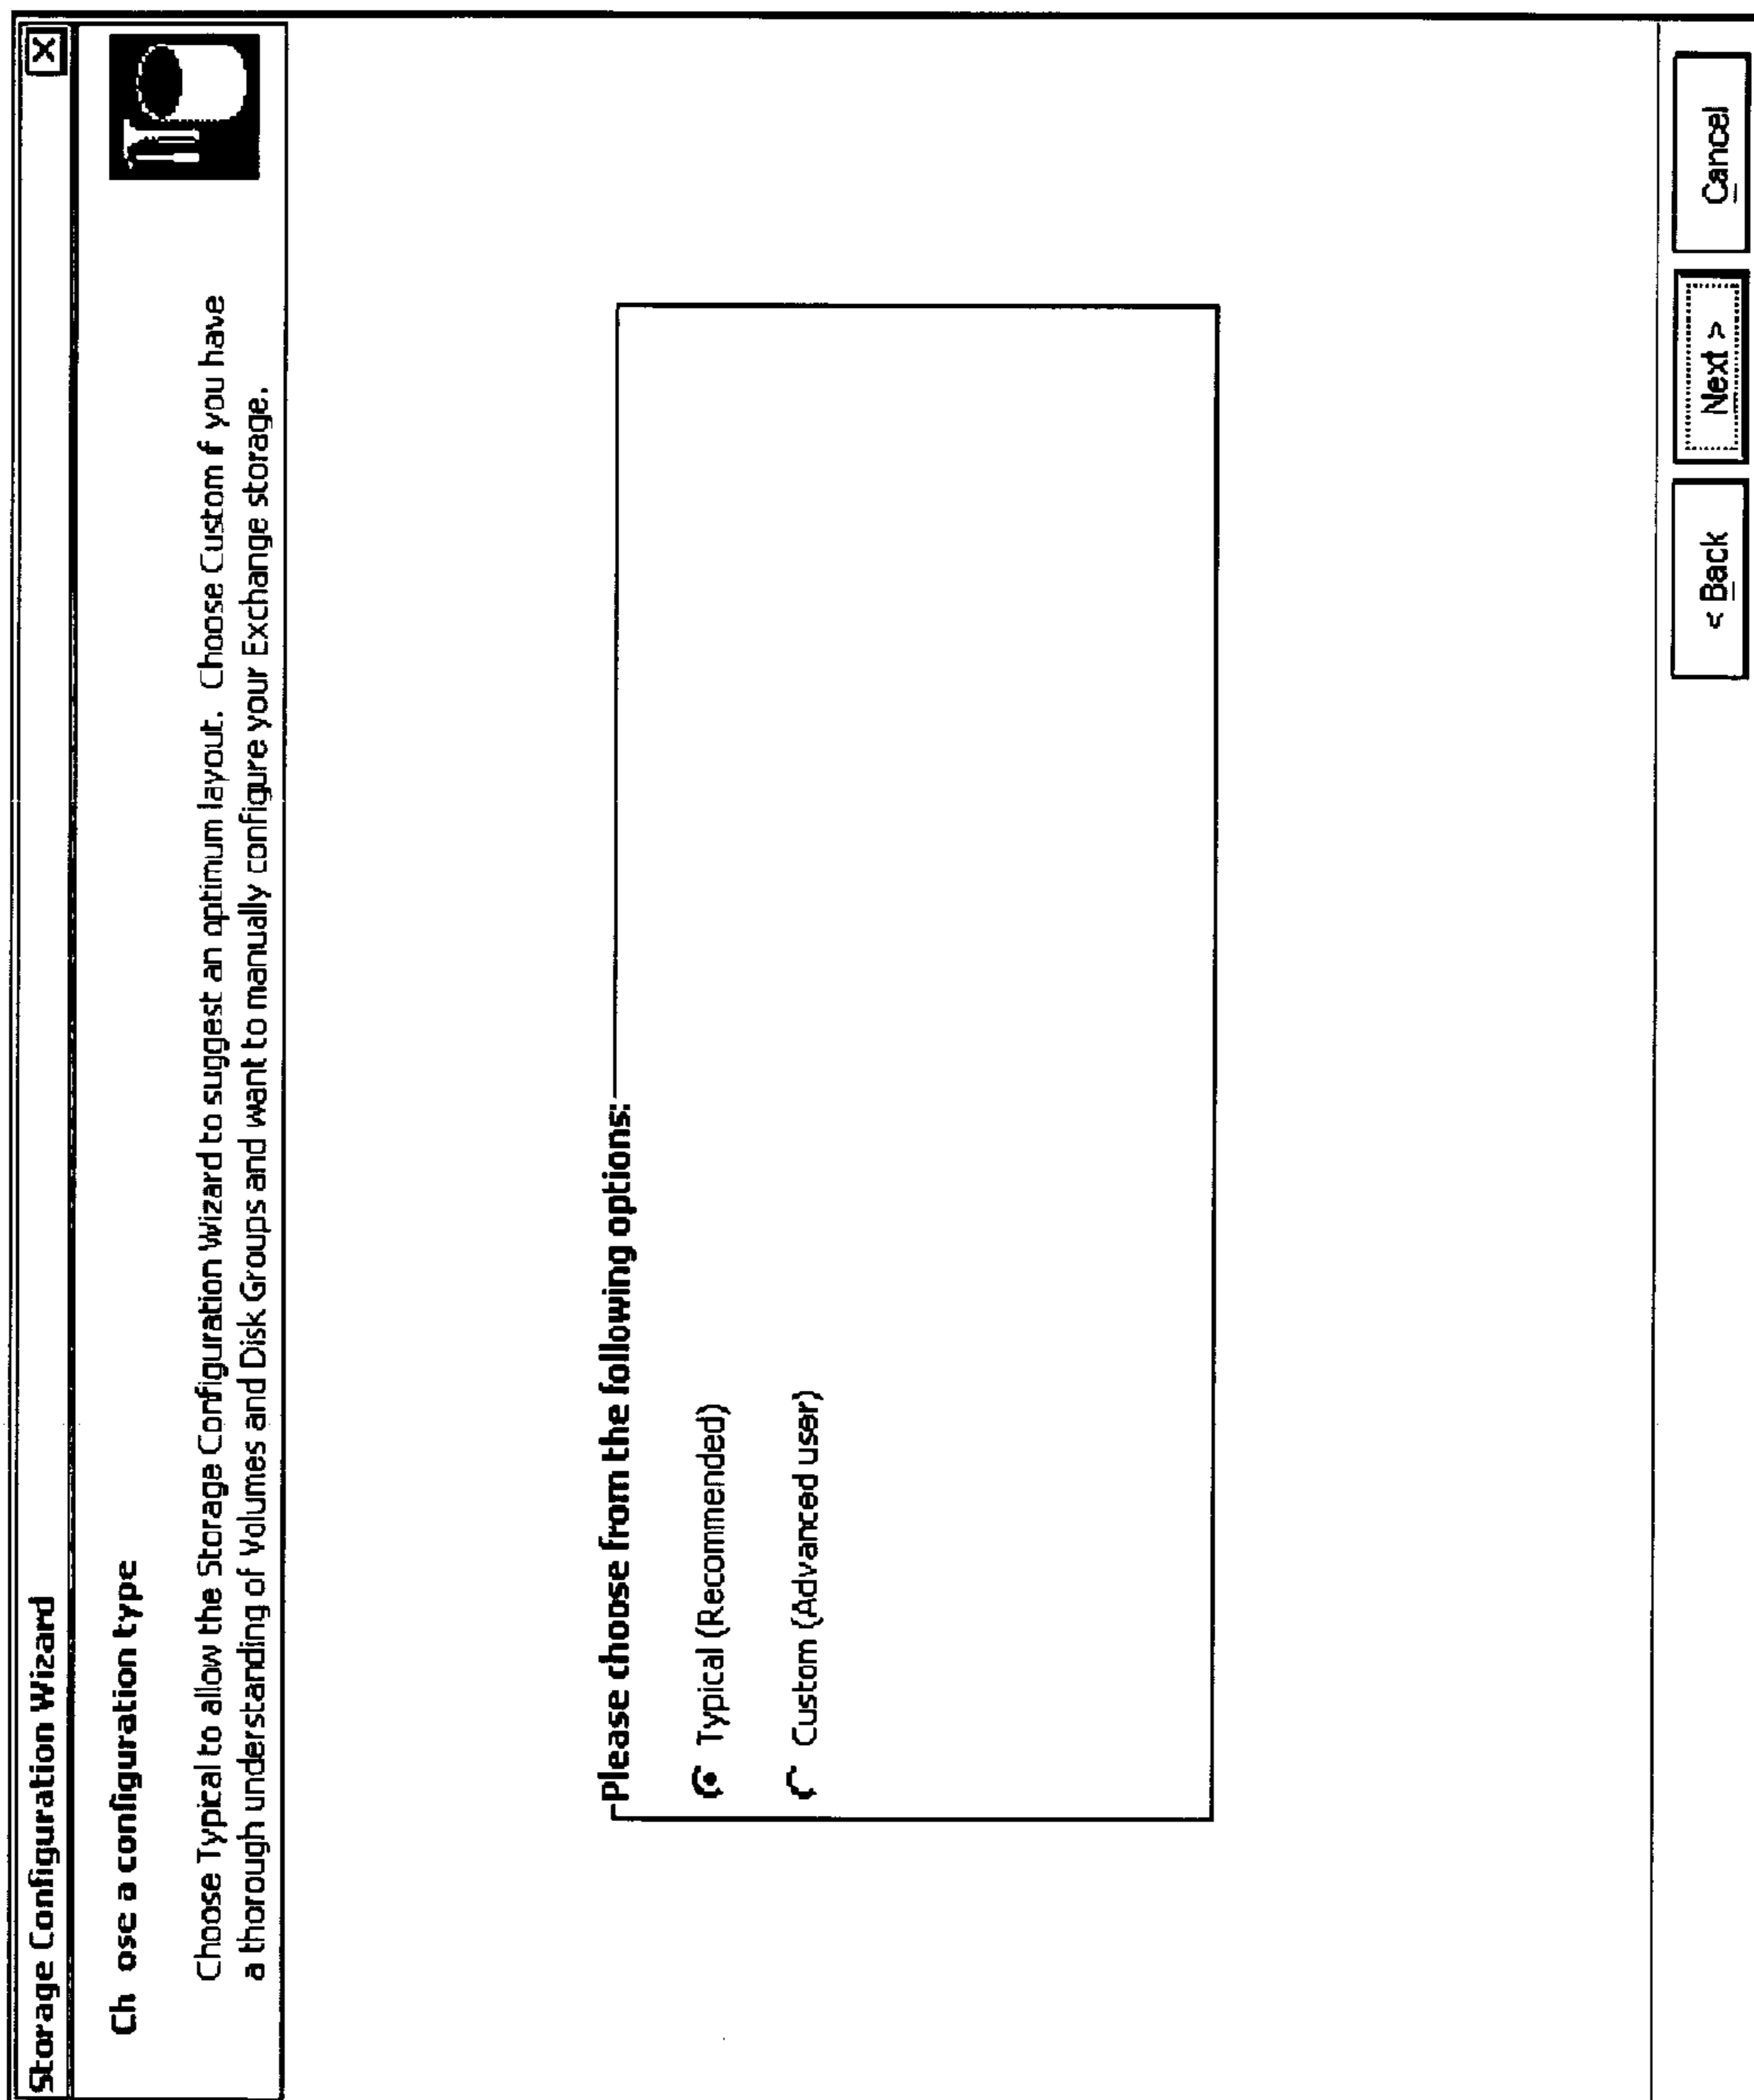


Fig. 4

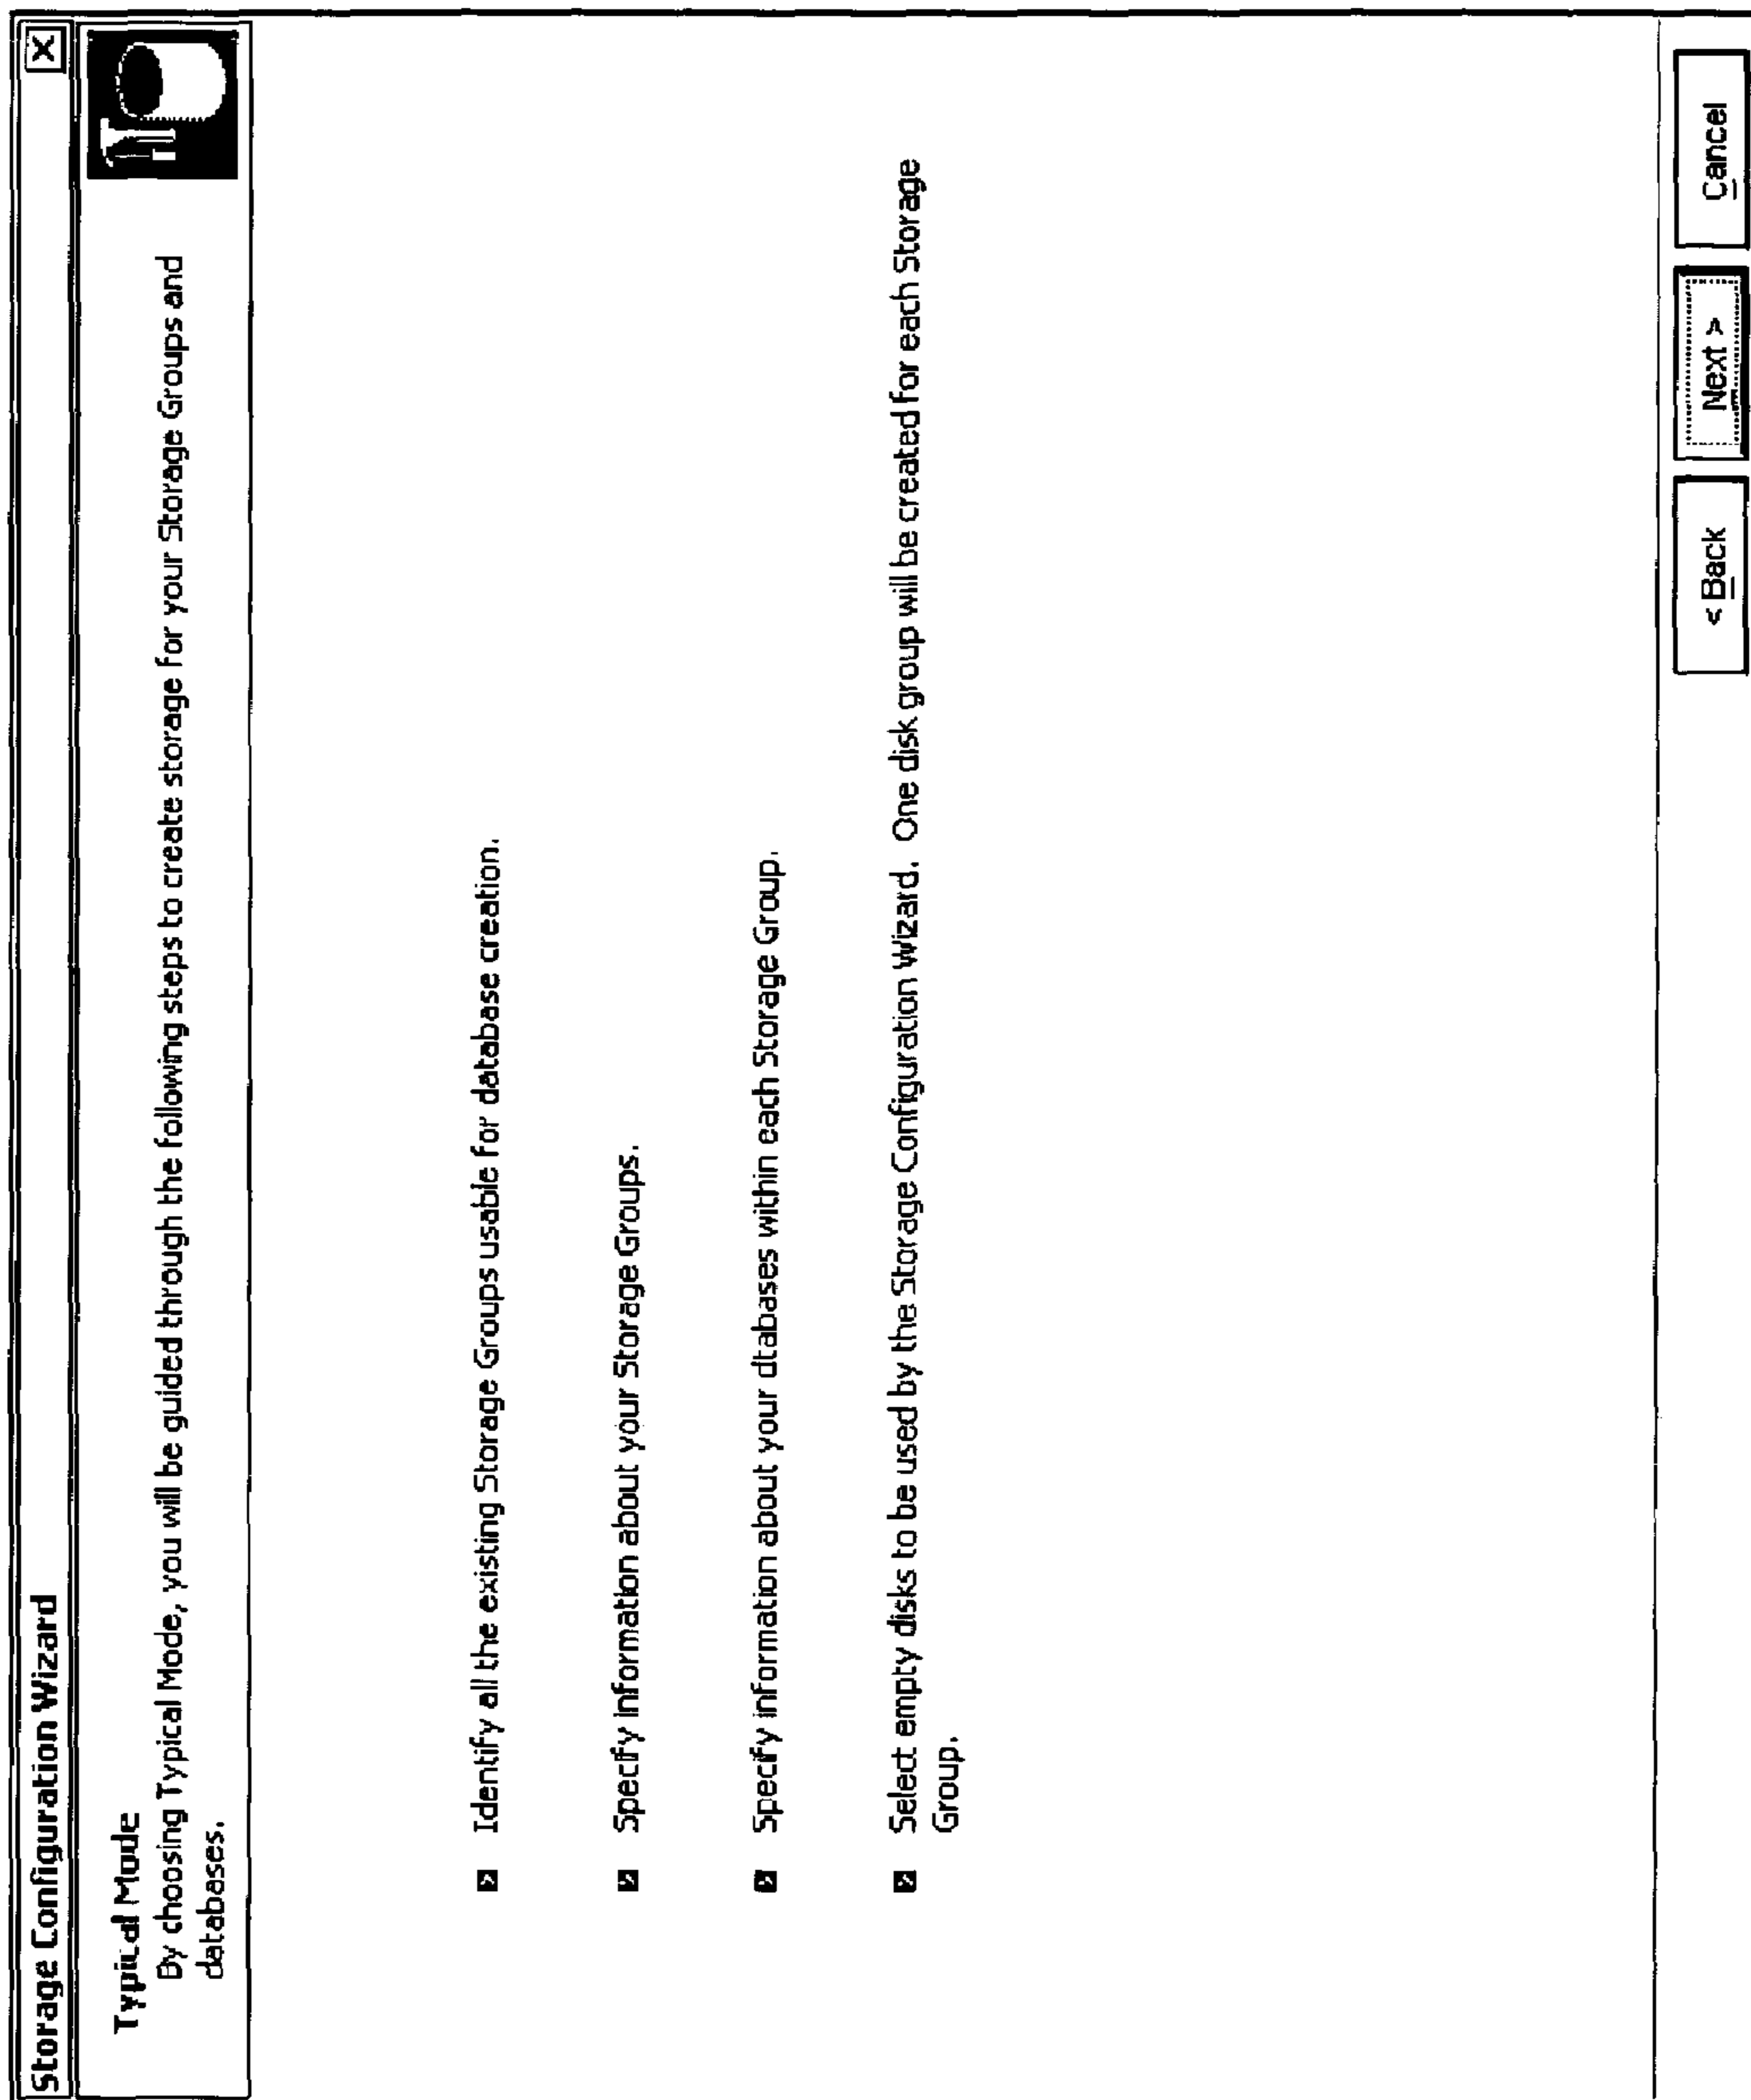


Fig. 5

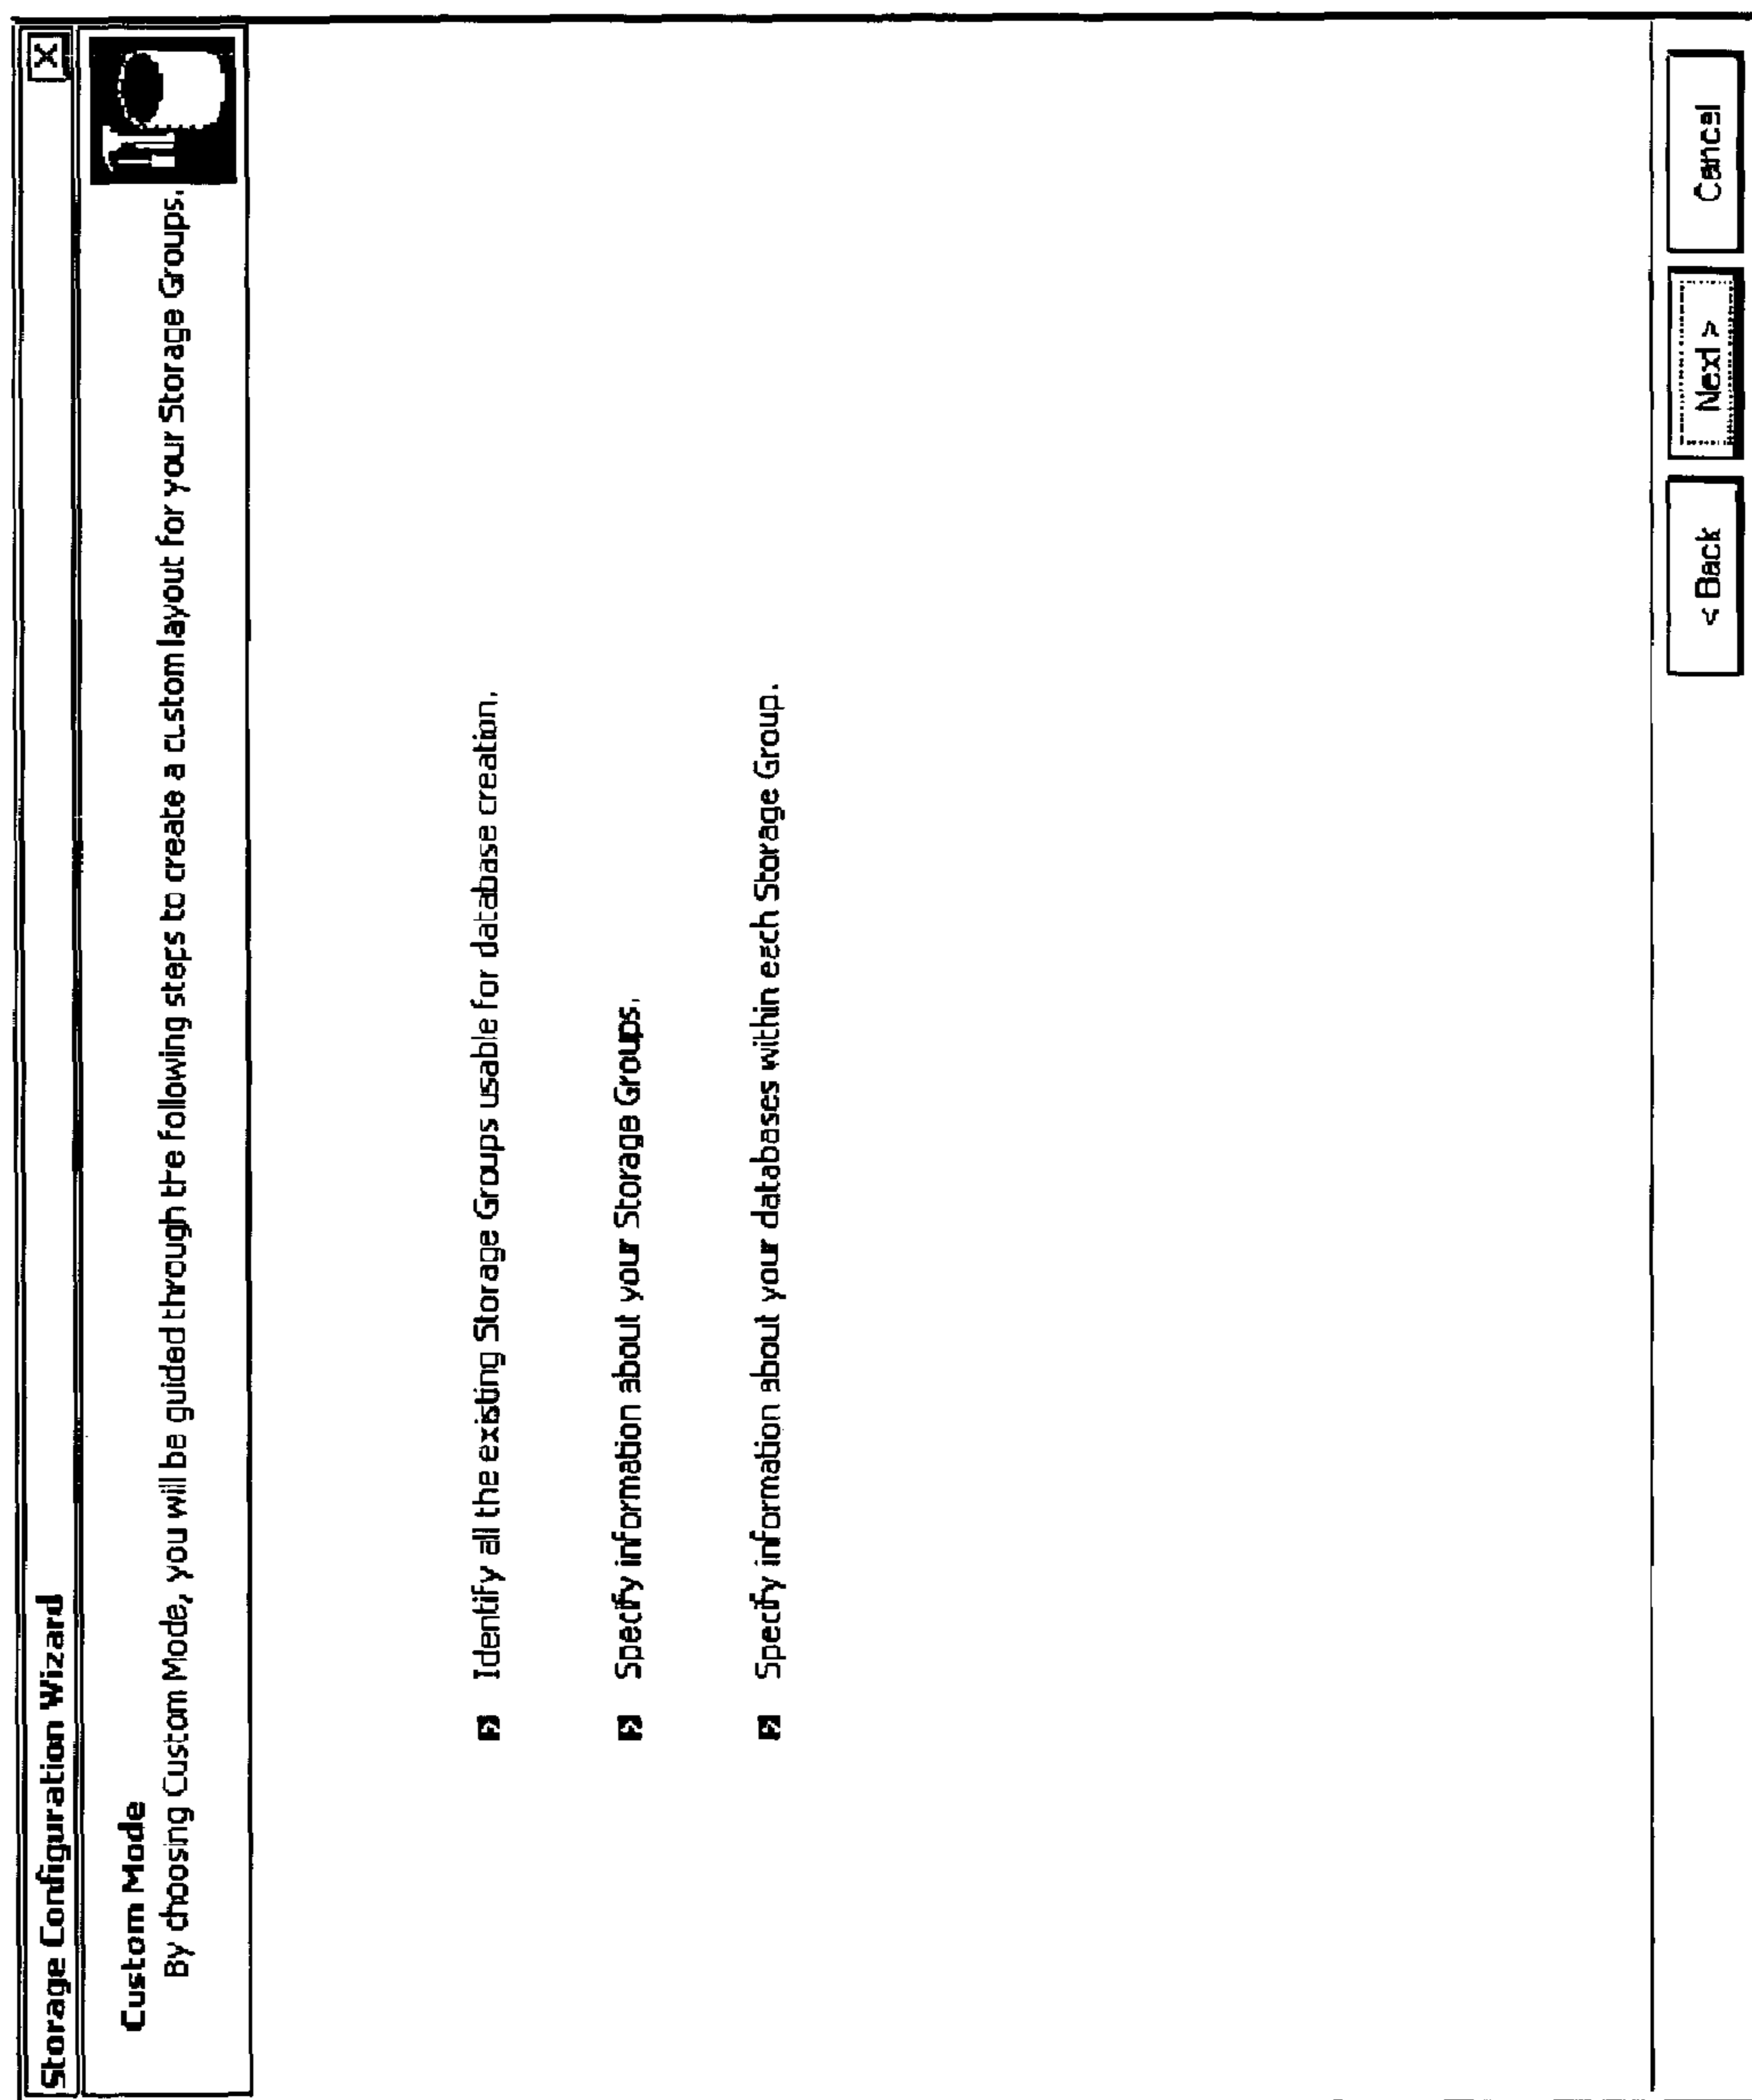


Fig. 6



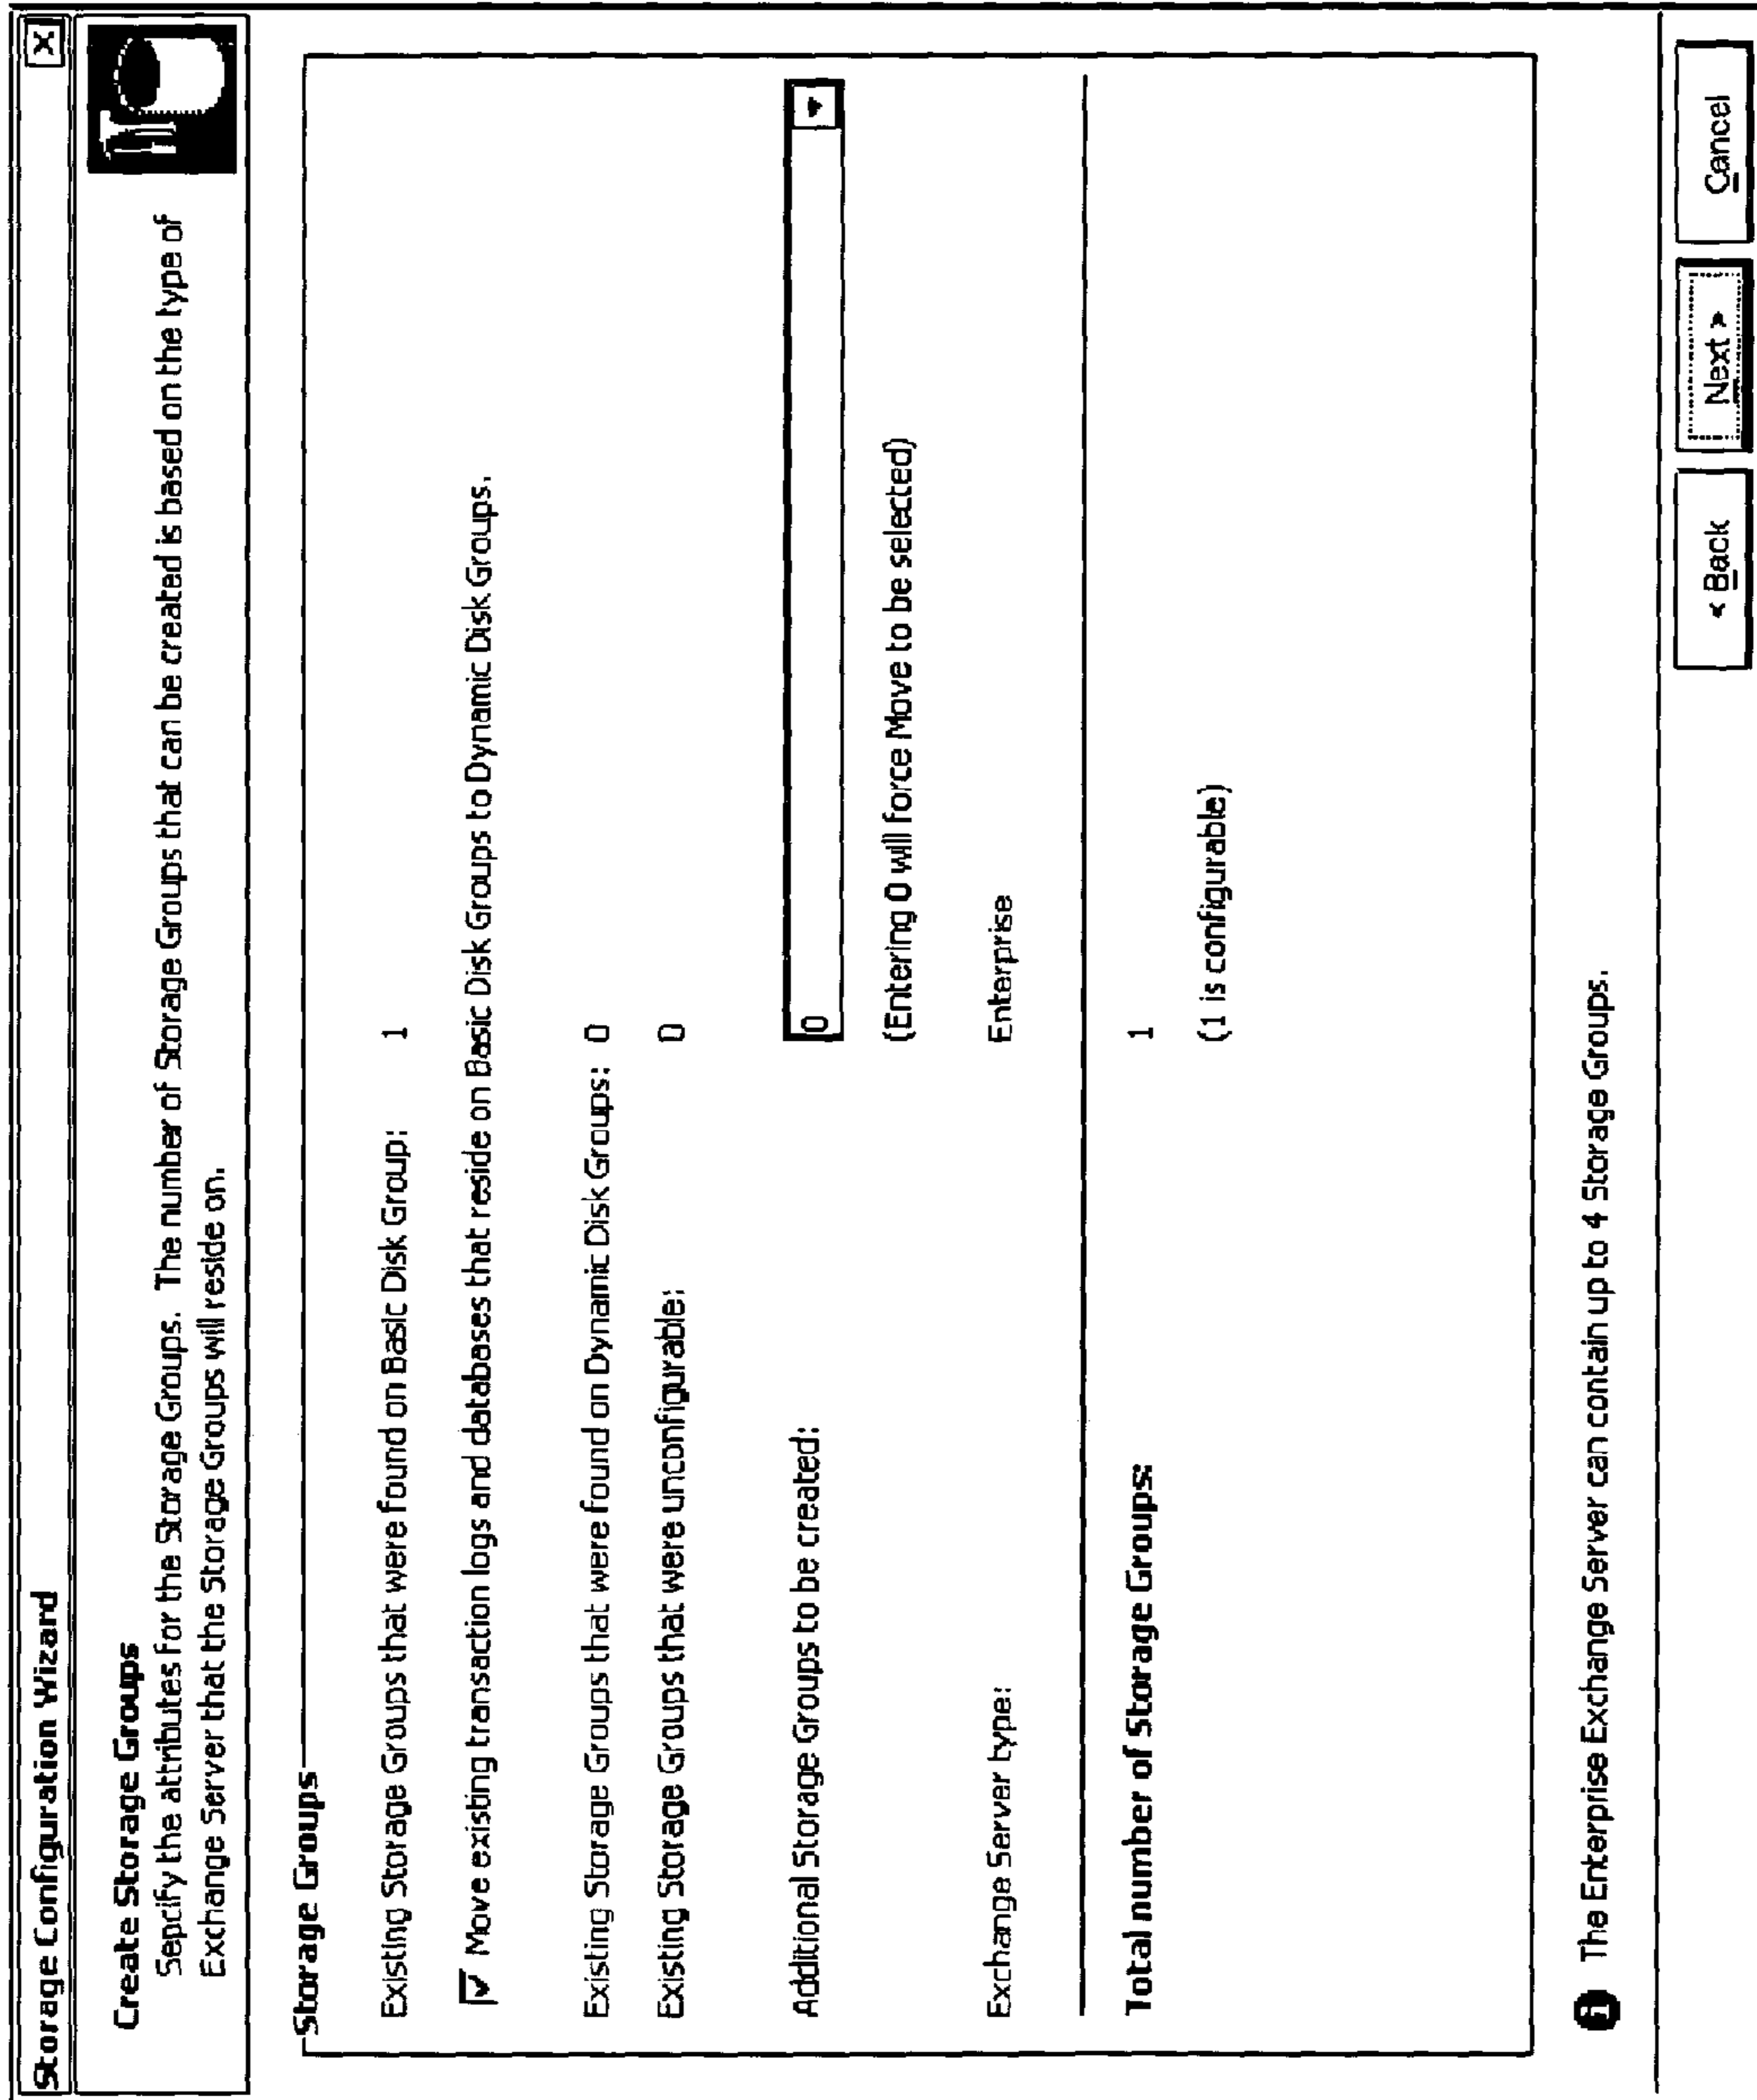


Fig. 7

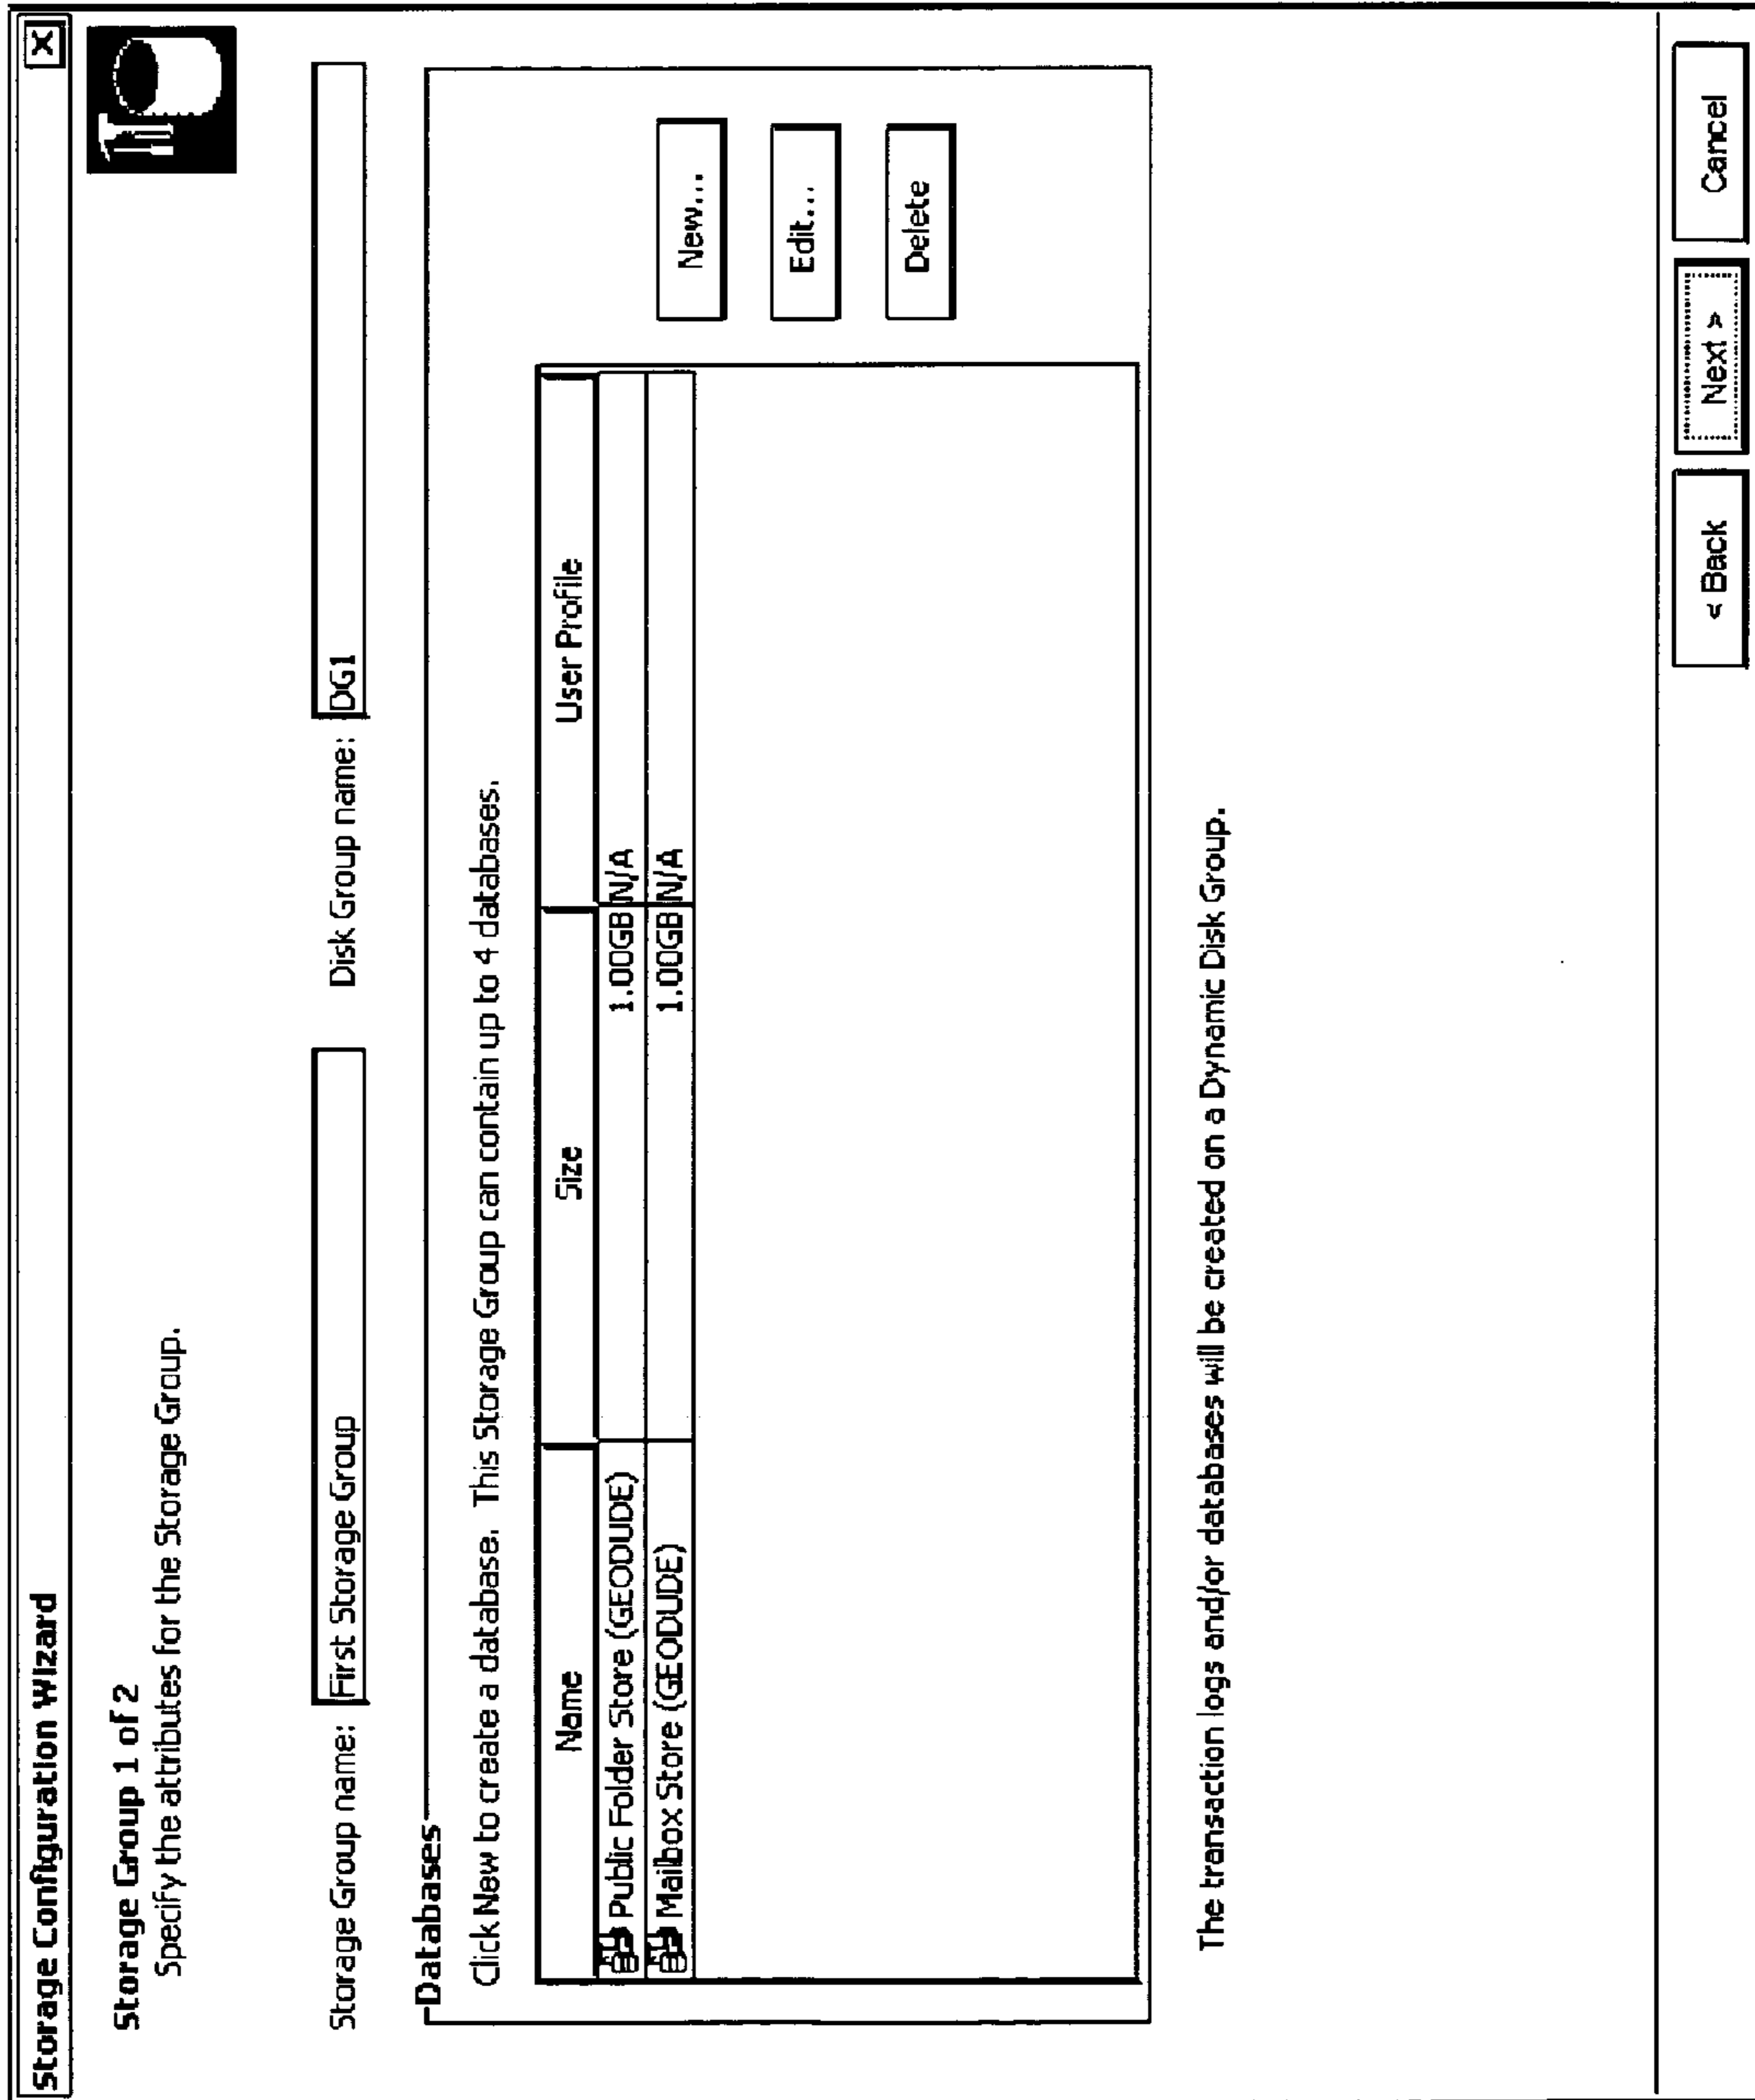
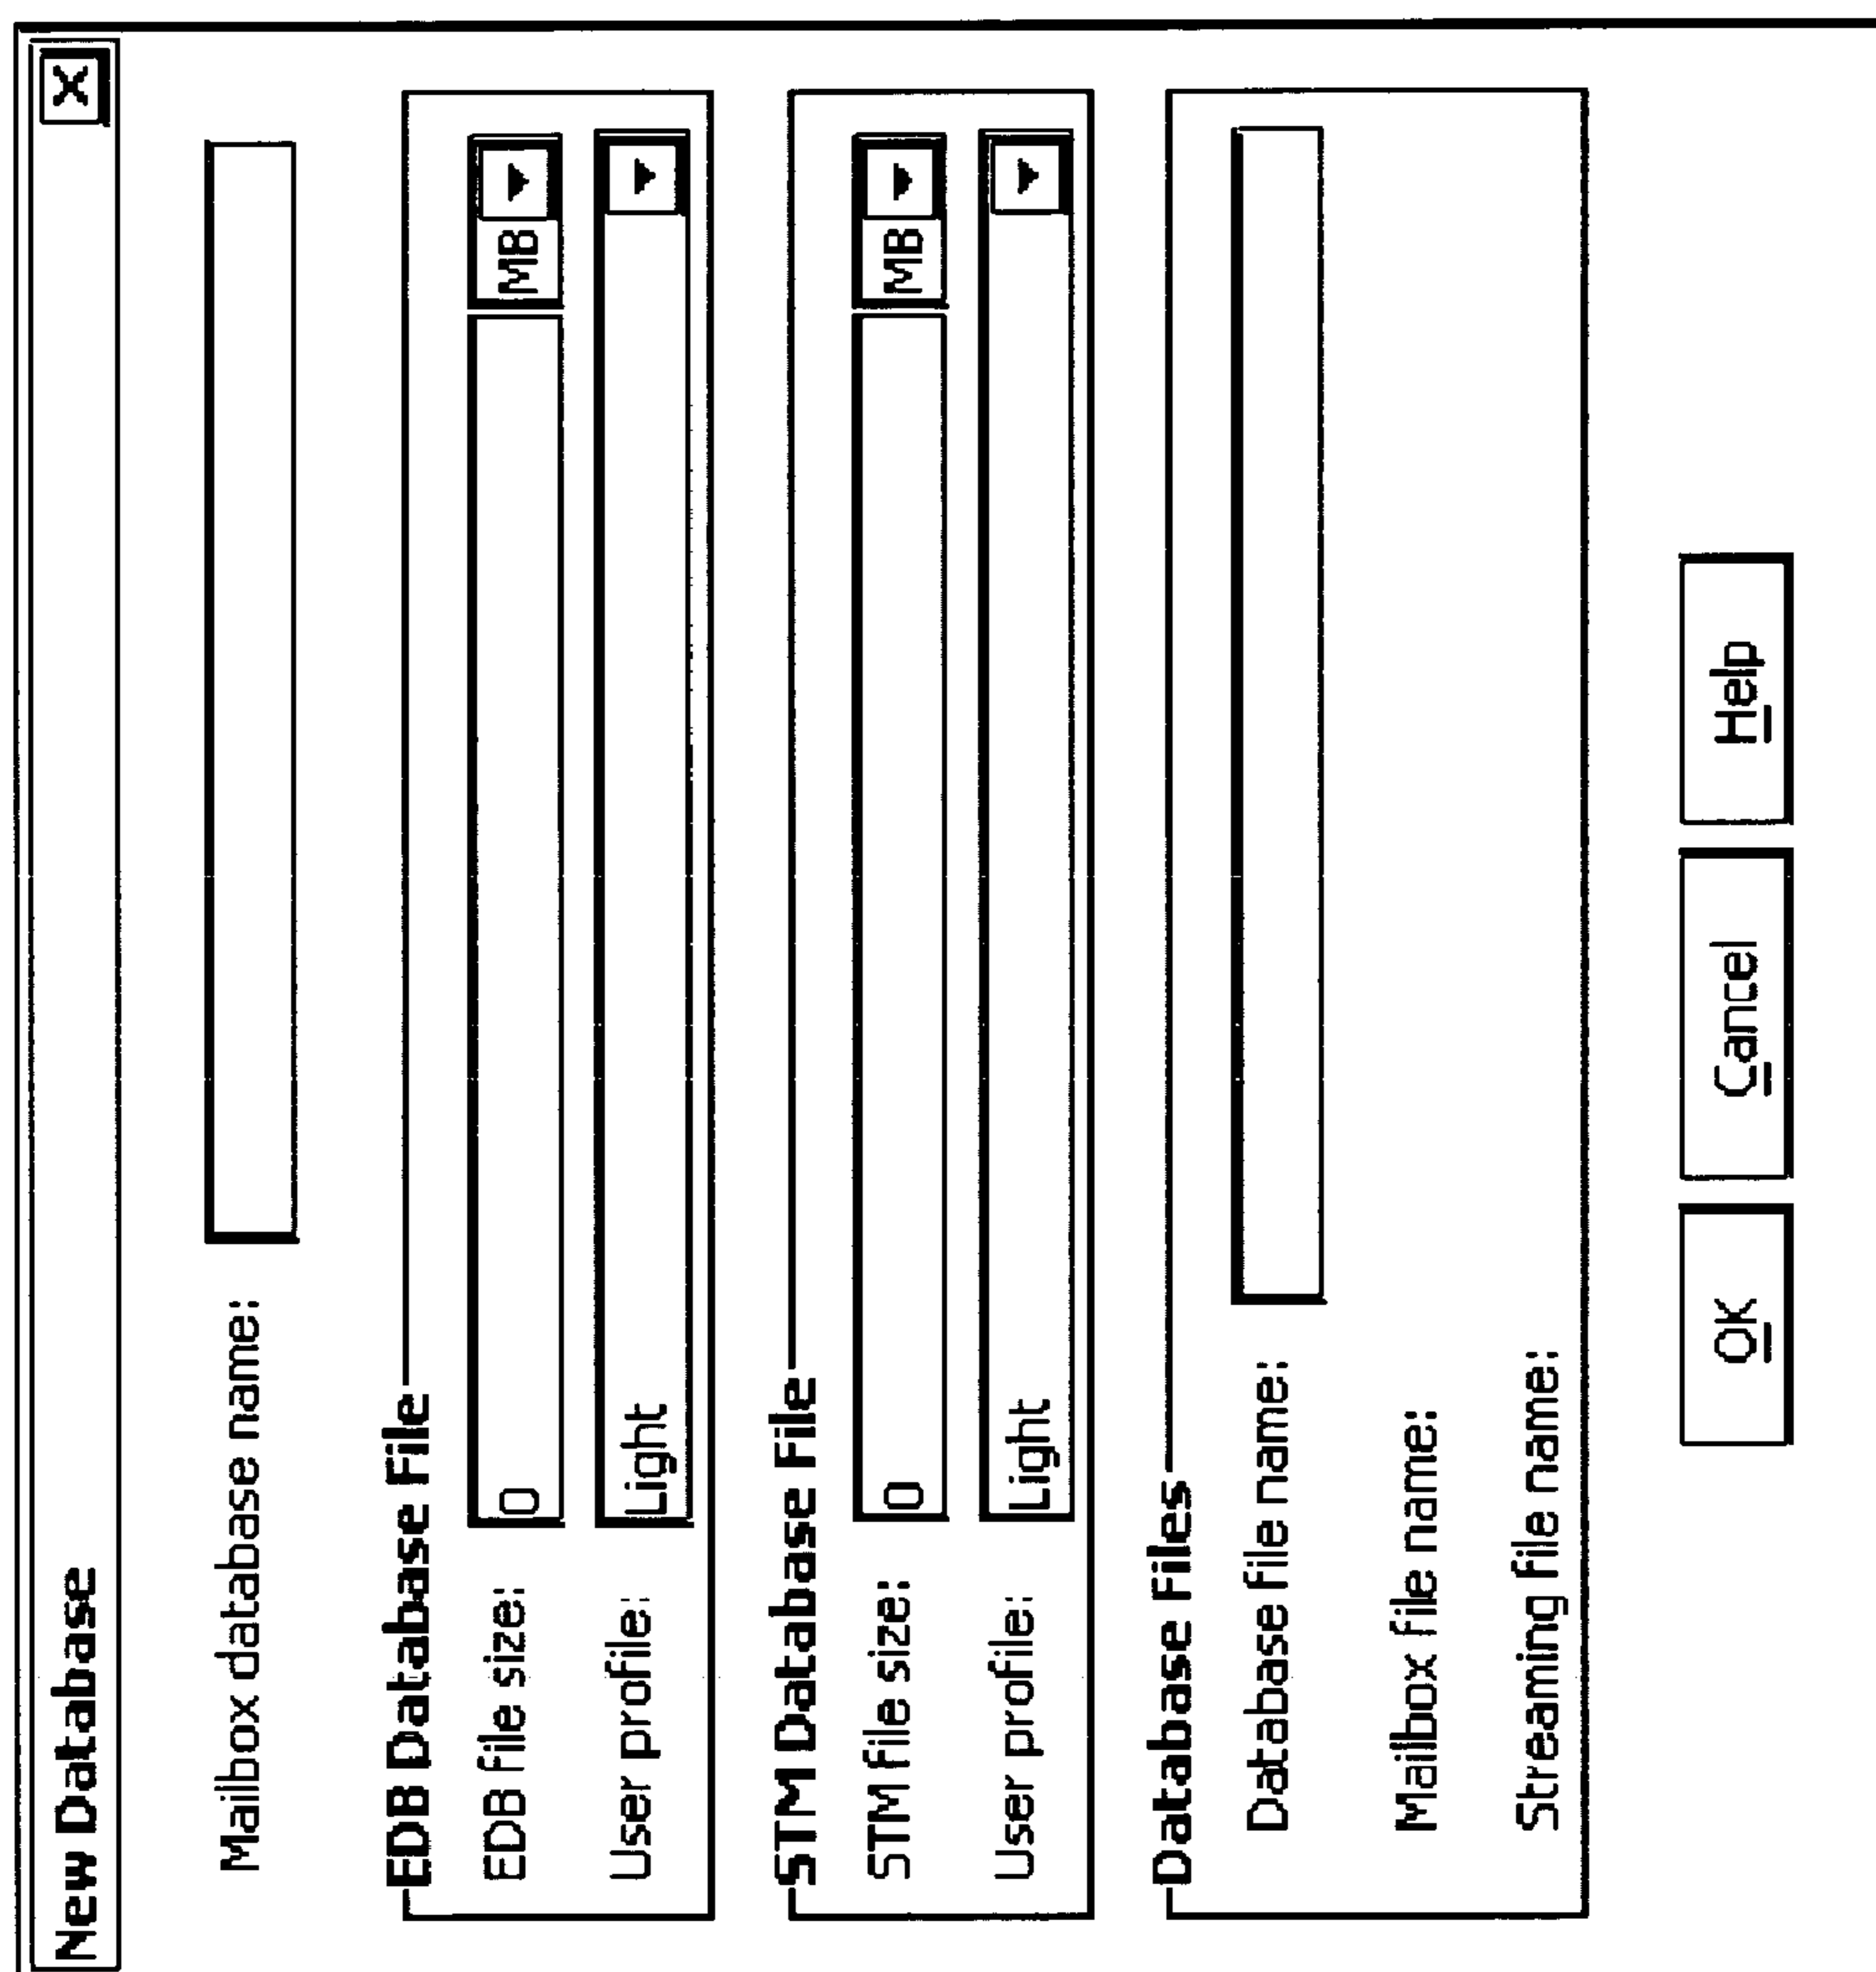


Fig. 8



The dialog box is titled "New Database" and contains the following fields and controls:

- Mailbox database name:** A text input field.
- EDB Database File:**
  - EDB file size:** A numeric input field with "0" and a unit selector dropdown set to "MB".
  - User profile:** A dropdown menu set to "Light".
- STM Database File:**
  - STM file size:** A numeric input field with "0" and a unit selector dropdown set to "MB".
  - User profile:** A dropdown menu set to "Light".
- Database Files:**
  - Database file name:** A text input field.
  - Mailbox file name:** A text input field.
  - Streaming file name:** A text input field.
- Buttons:** "OK", "Cancel", and "Help" buttons are located at the bottom right.

Fig. 9

**Storage Configuration Wizard**

**Storage Group 1 of 2**  
Specify the attributes for the Storage Group.

Storage Group name:  Disk Group name:

**Databases**  
Click New to create a database. This Storage Group can contain up to 4 databases (9 Exchange objects).  
Select a database file and click "Set Volume" to assign a volume.

Transaction Log/D...	File Name	Assigned Volume	Size	User Profile
First Storage G...	Multiple Files		0.50GB	N/A
Public Folder S...	Public Folder Store ...		0.50GB	N/A
Public Folder S...	Public Folder Store ...		0.50GB	N/A
Mailbox Store (...)	Mailbox Store (GE...		0.50GB	N/A
Mailbox Store (...)	Mailbox Store (GE...		0.50GB	N/A

Set Volume... Hot Spare... New... Edit... Delete

The transaction logs and/or databases will be created on a Dynamic Disk Group.

< Back Next > Cancel

Fig. 10

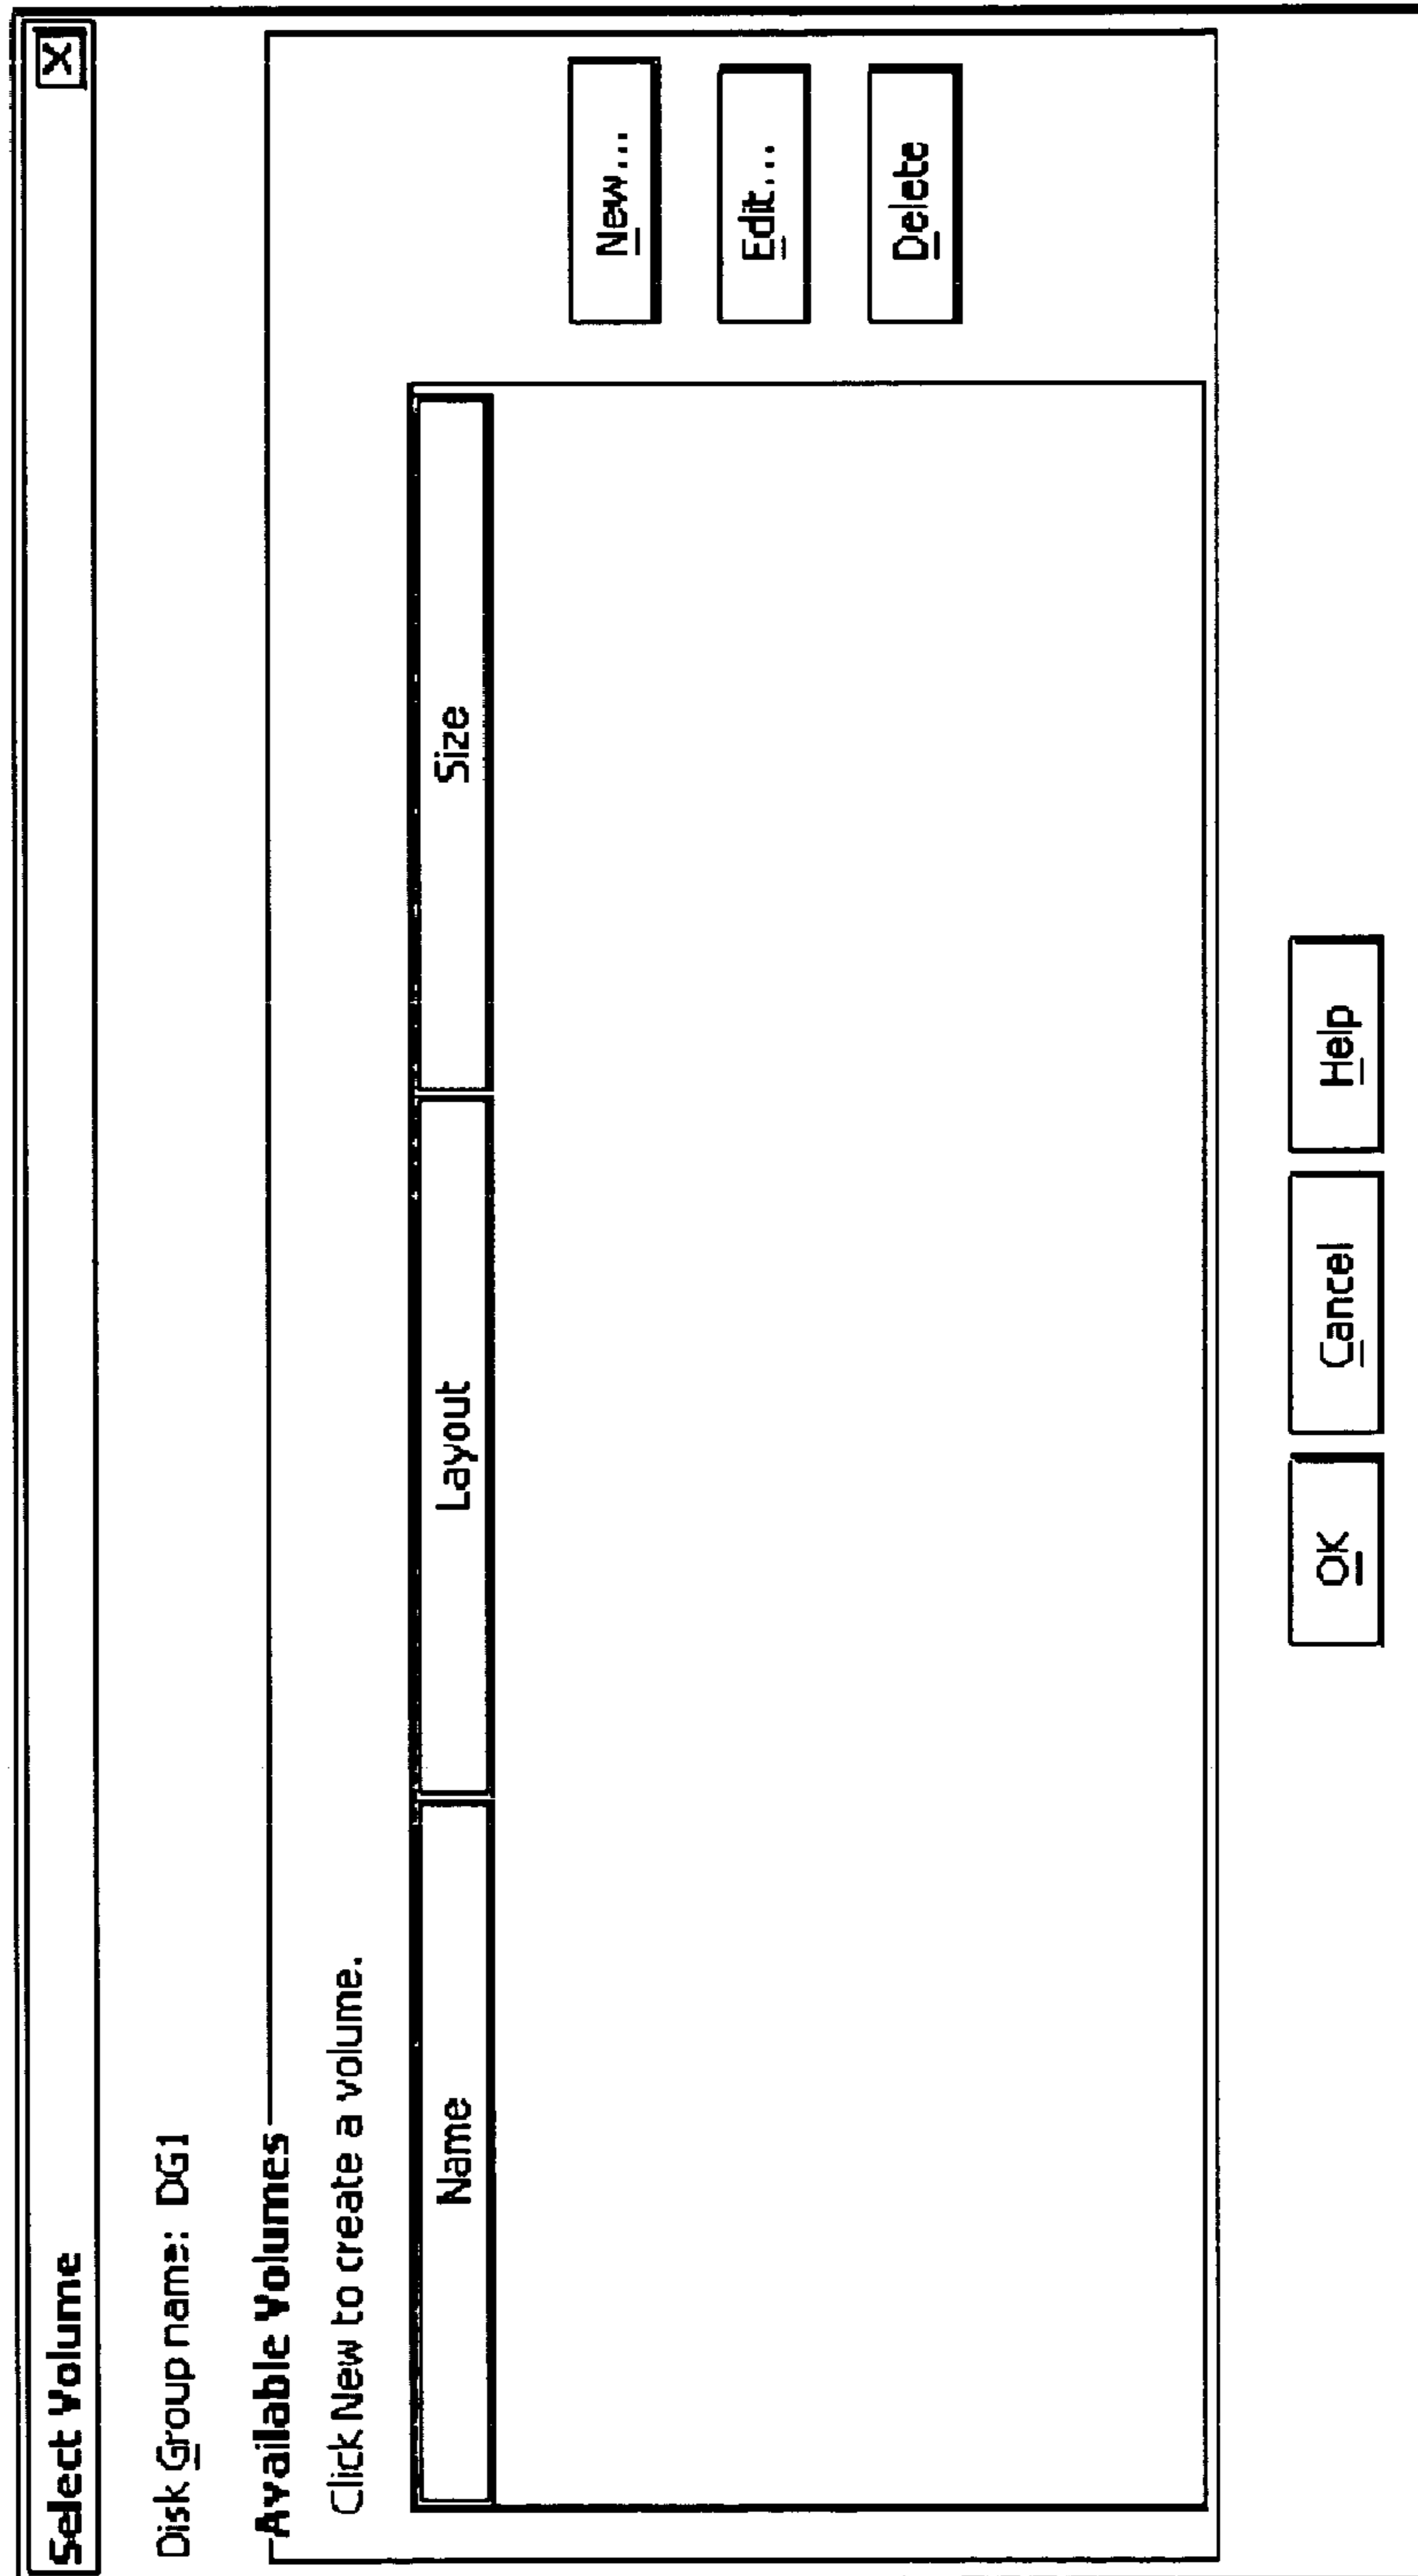


Fig. 11

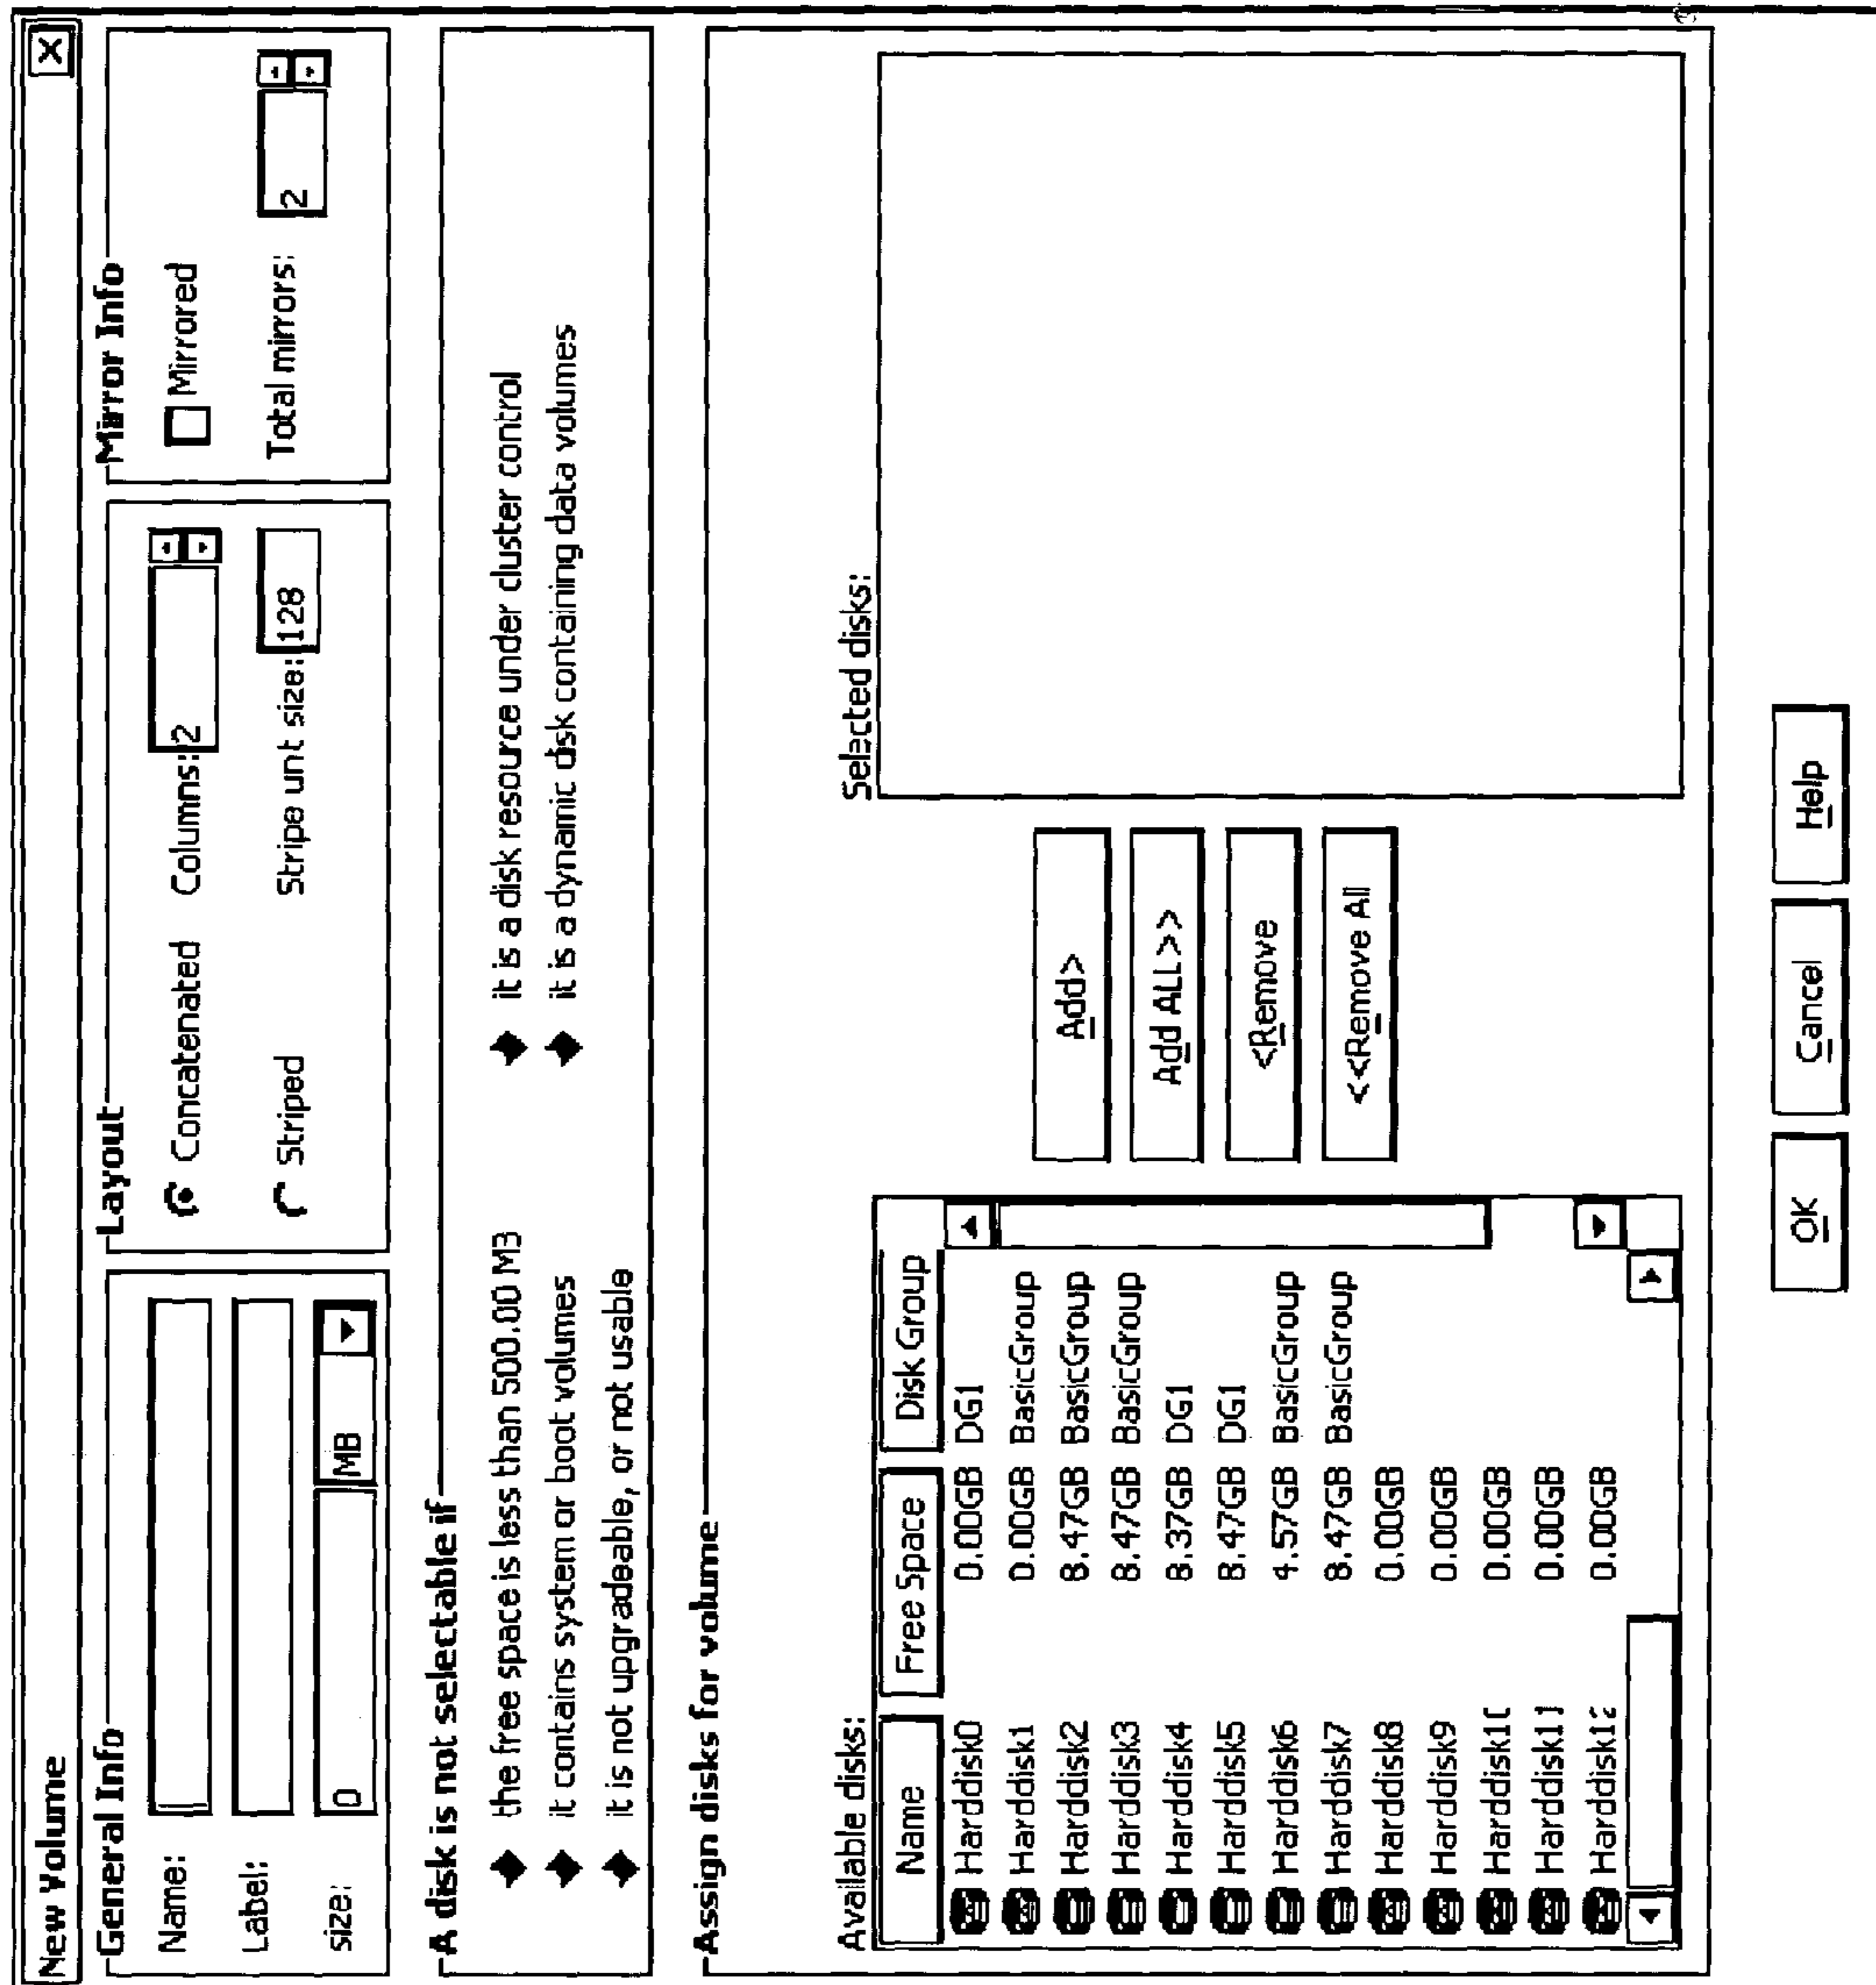


Fig. 12



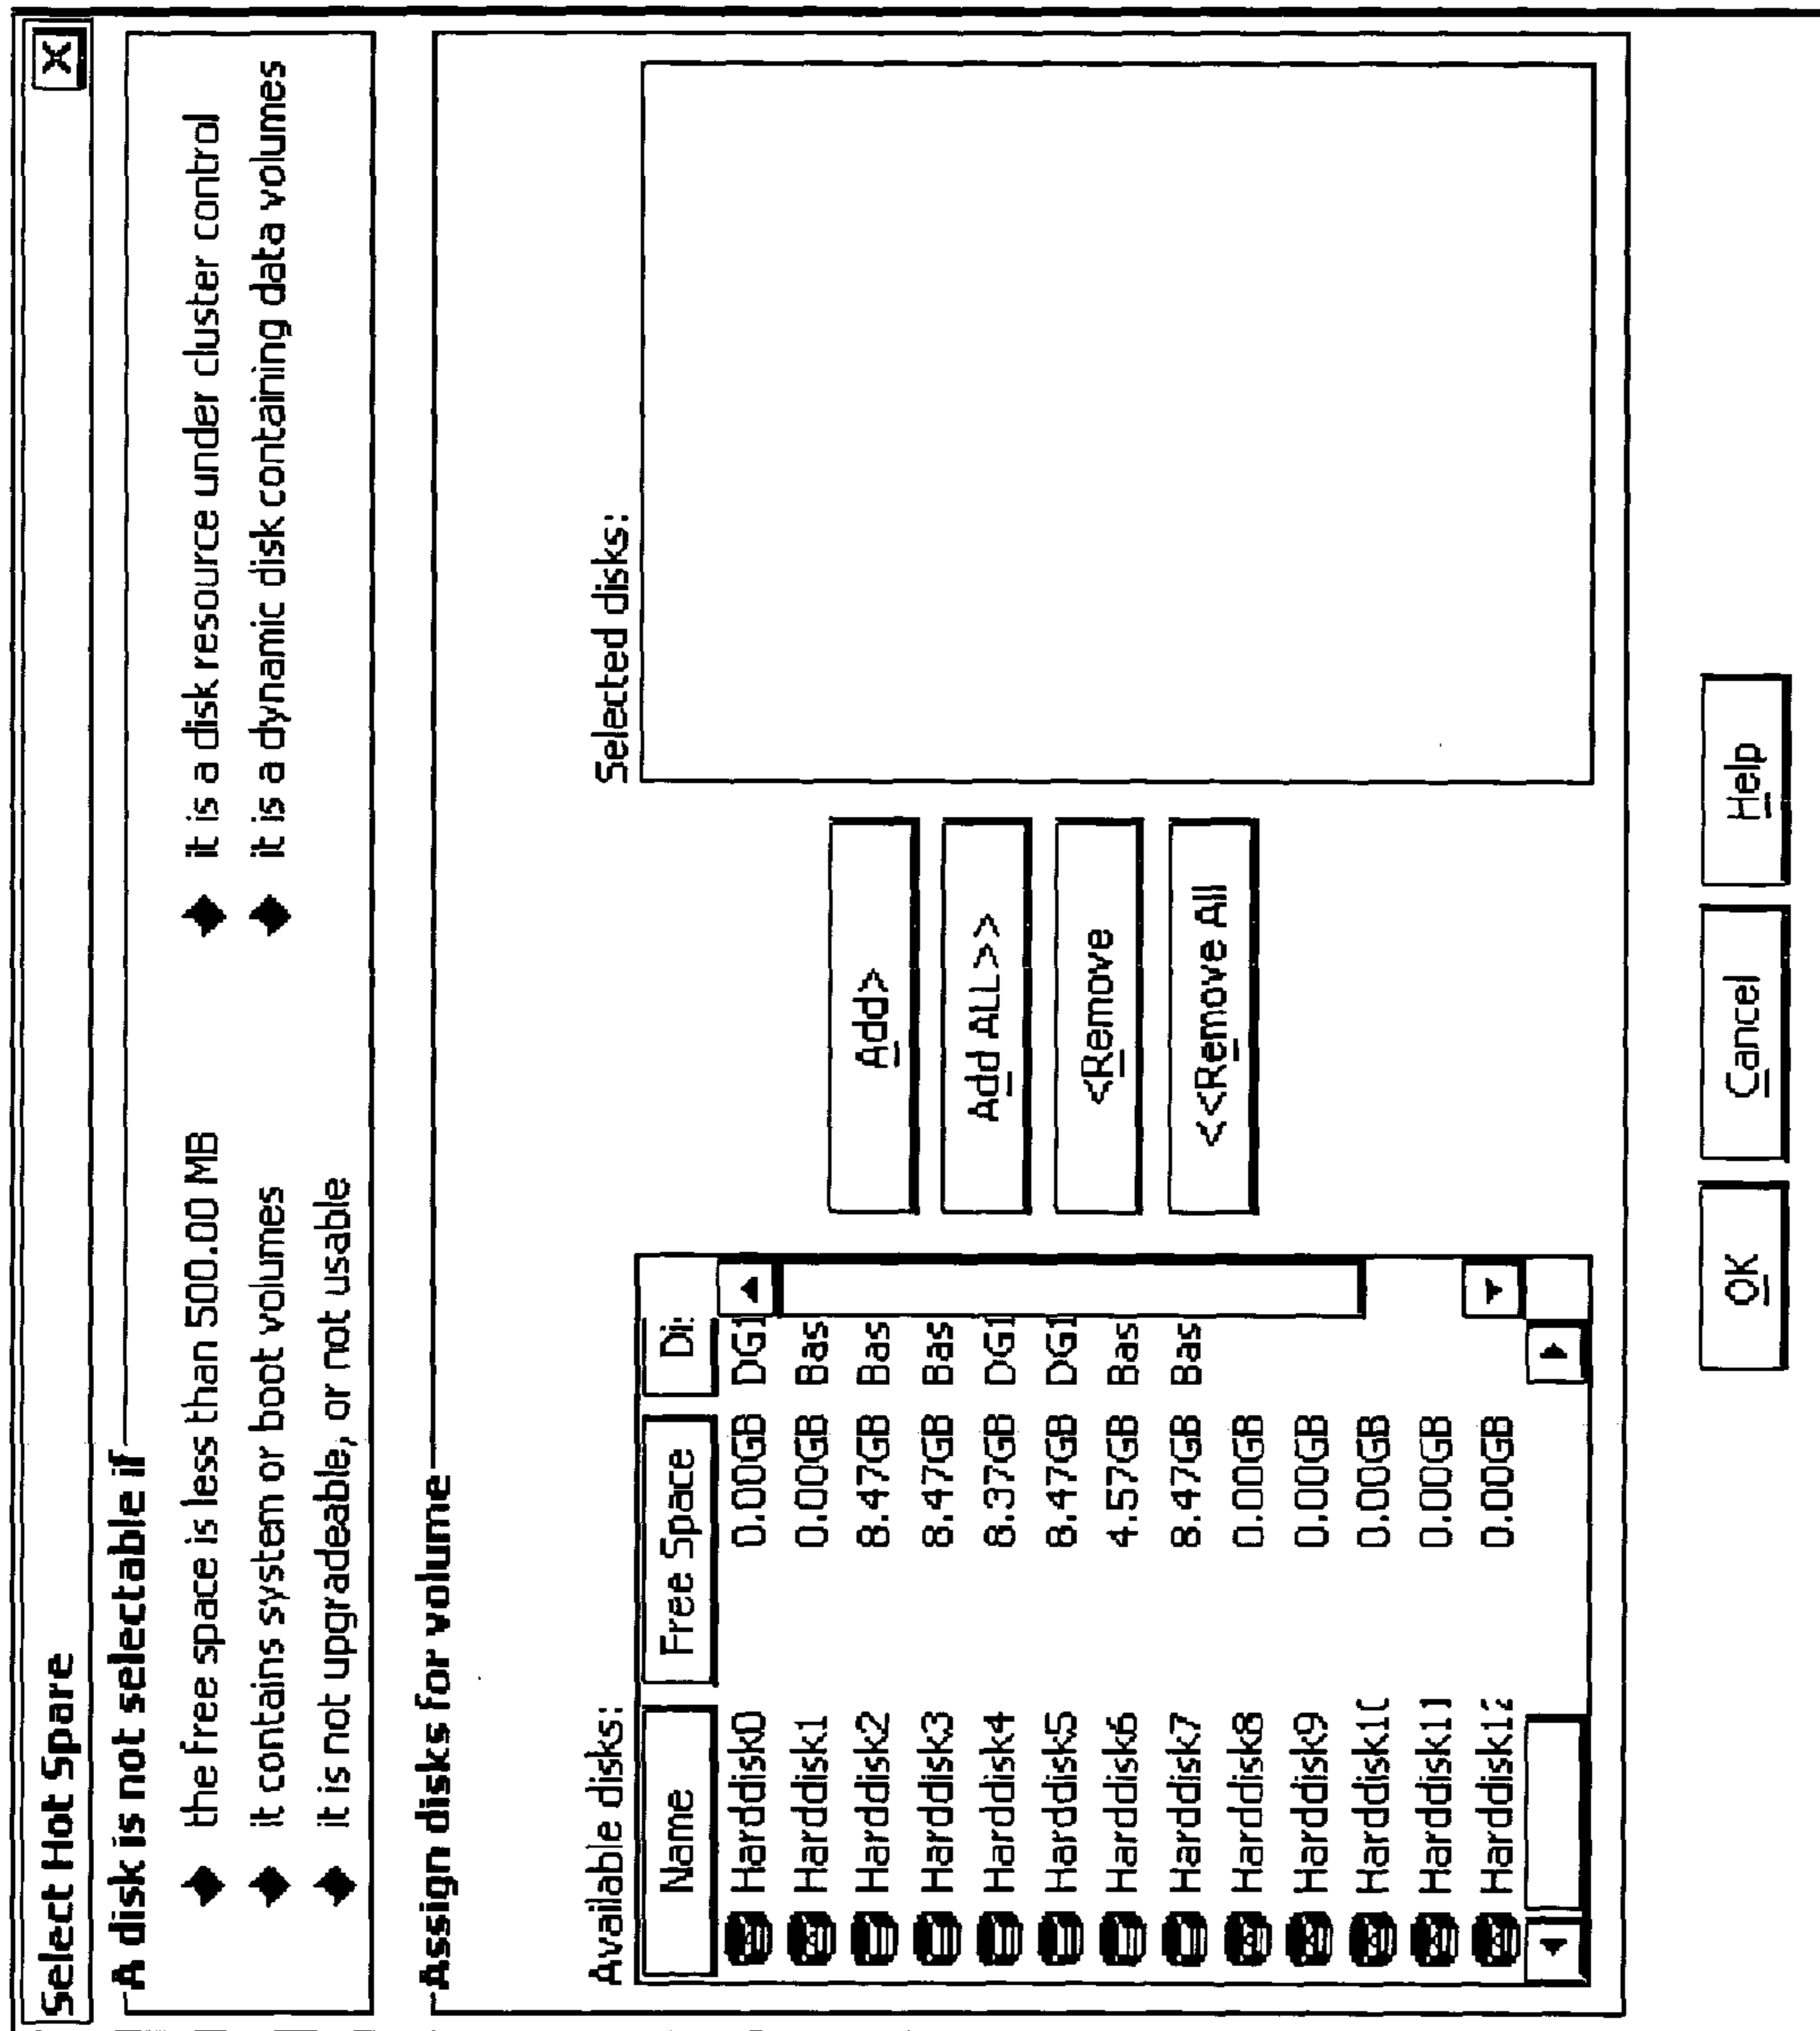


Fig. 13

The dialog box titled "Edit STM Database File" contains the following fields and controls:

- Mailbox database name:
- STM file size:  MB
- User profile:
- Database Files** (grouped in a dashed box):
  - Database file name:
  - Mailbox file name: newdb.edb
  - Streaming file name: newdb.stm
- Buttons:

Fig. 14

The dialog box titled "Edit Transaction Logs" contains the following fields and controls:

- Transaction Logs:
- Transaction logs size:  MB
- User profile:
- Buttons:

Fig. 15

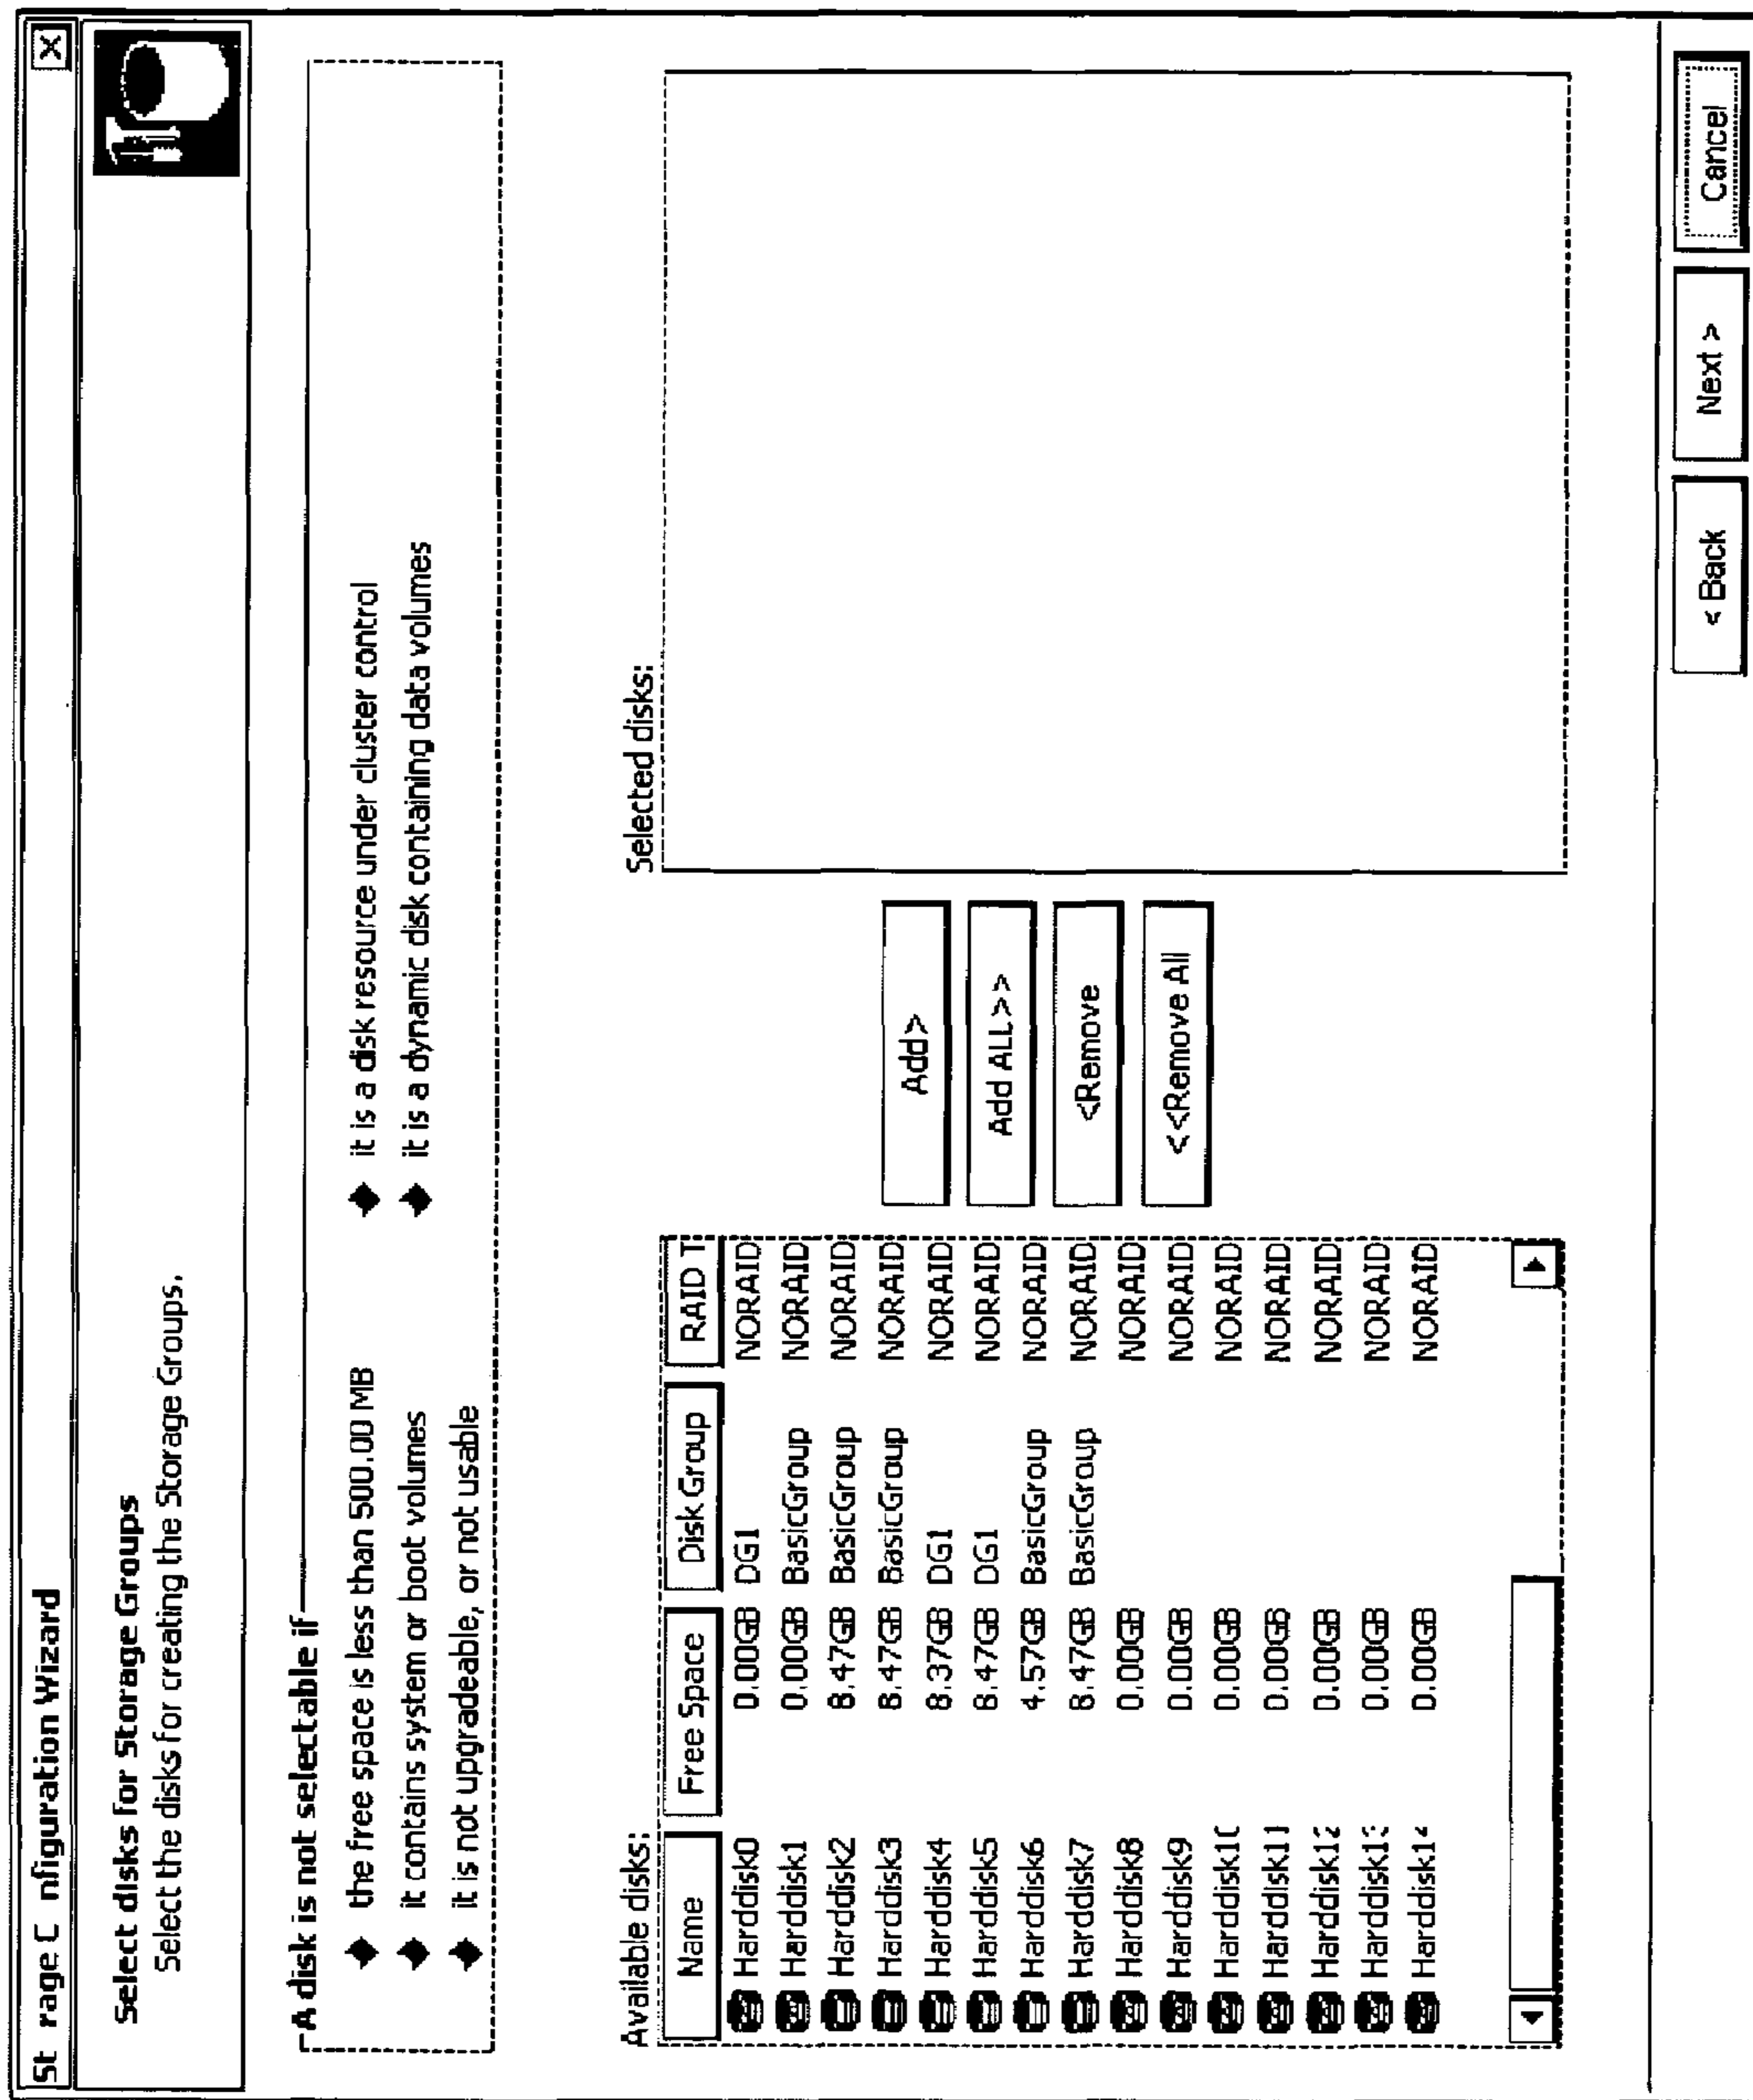


Fig. 16

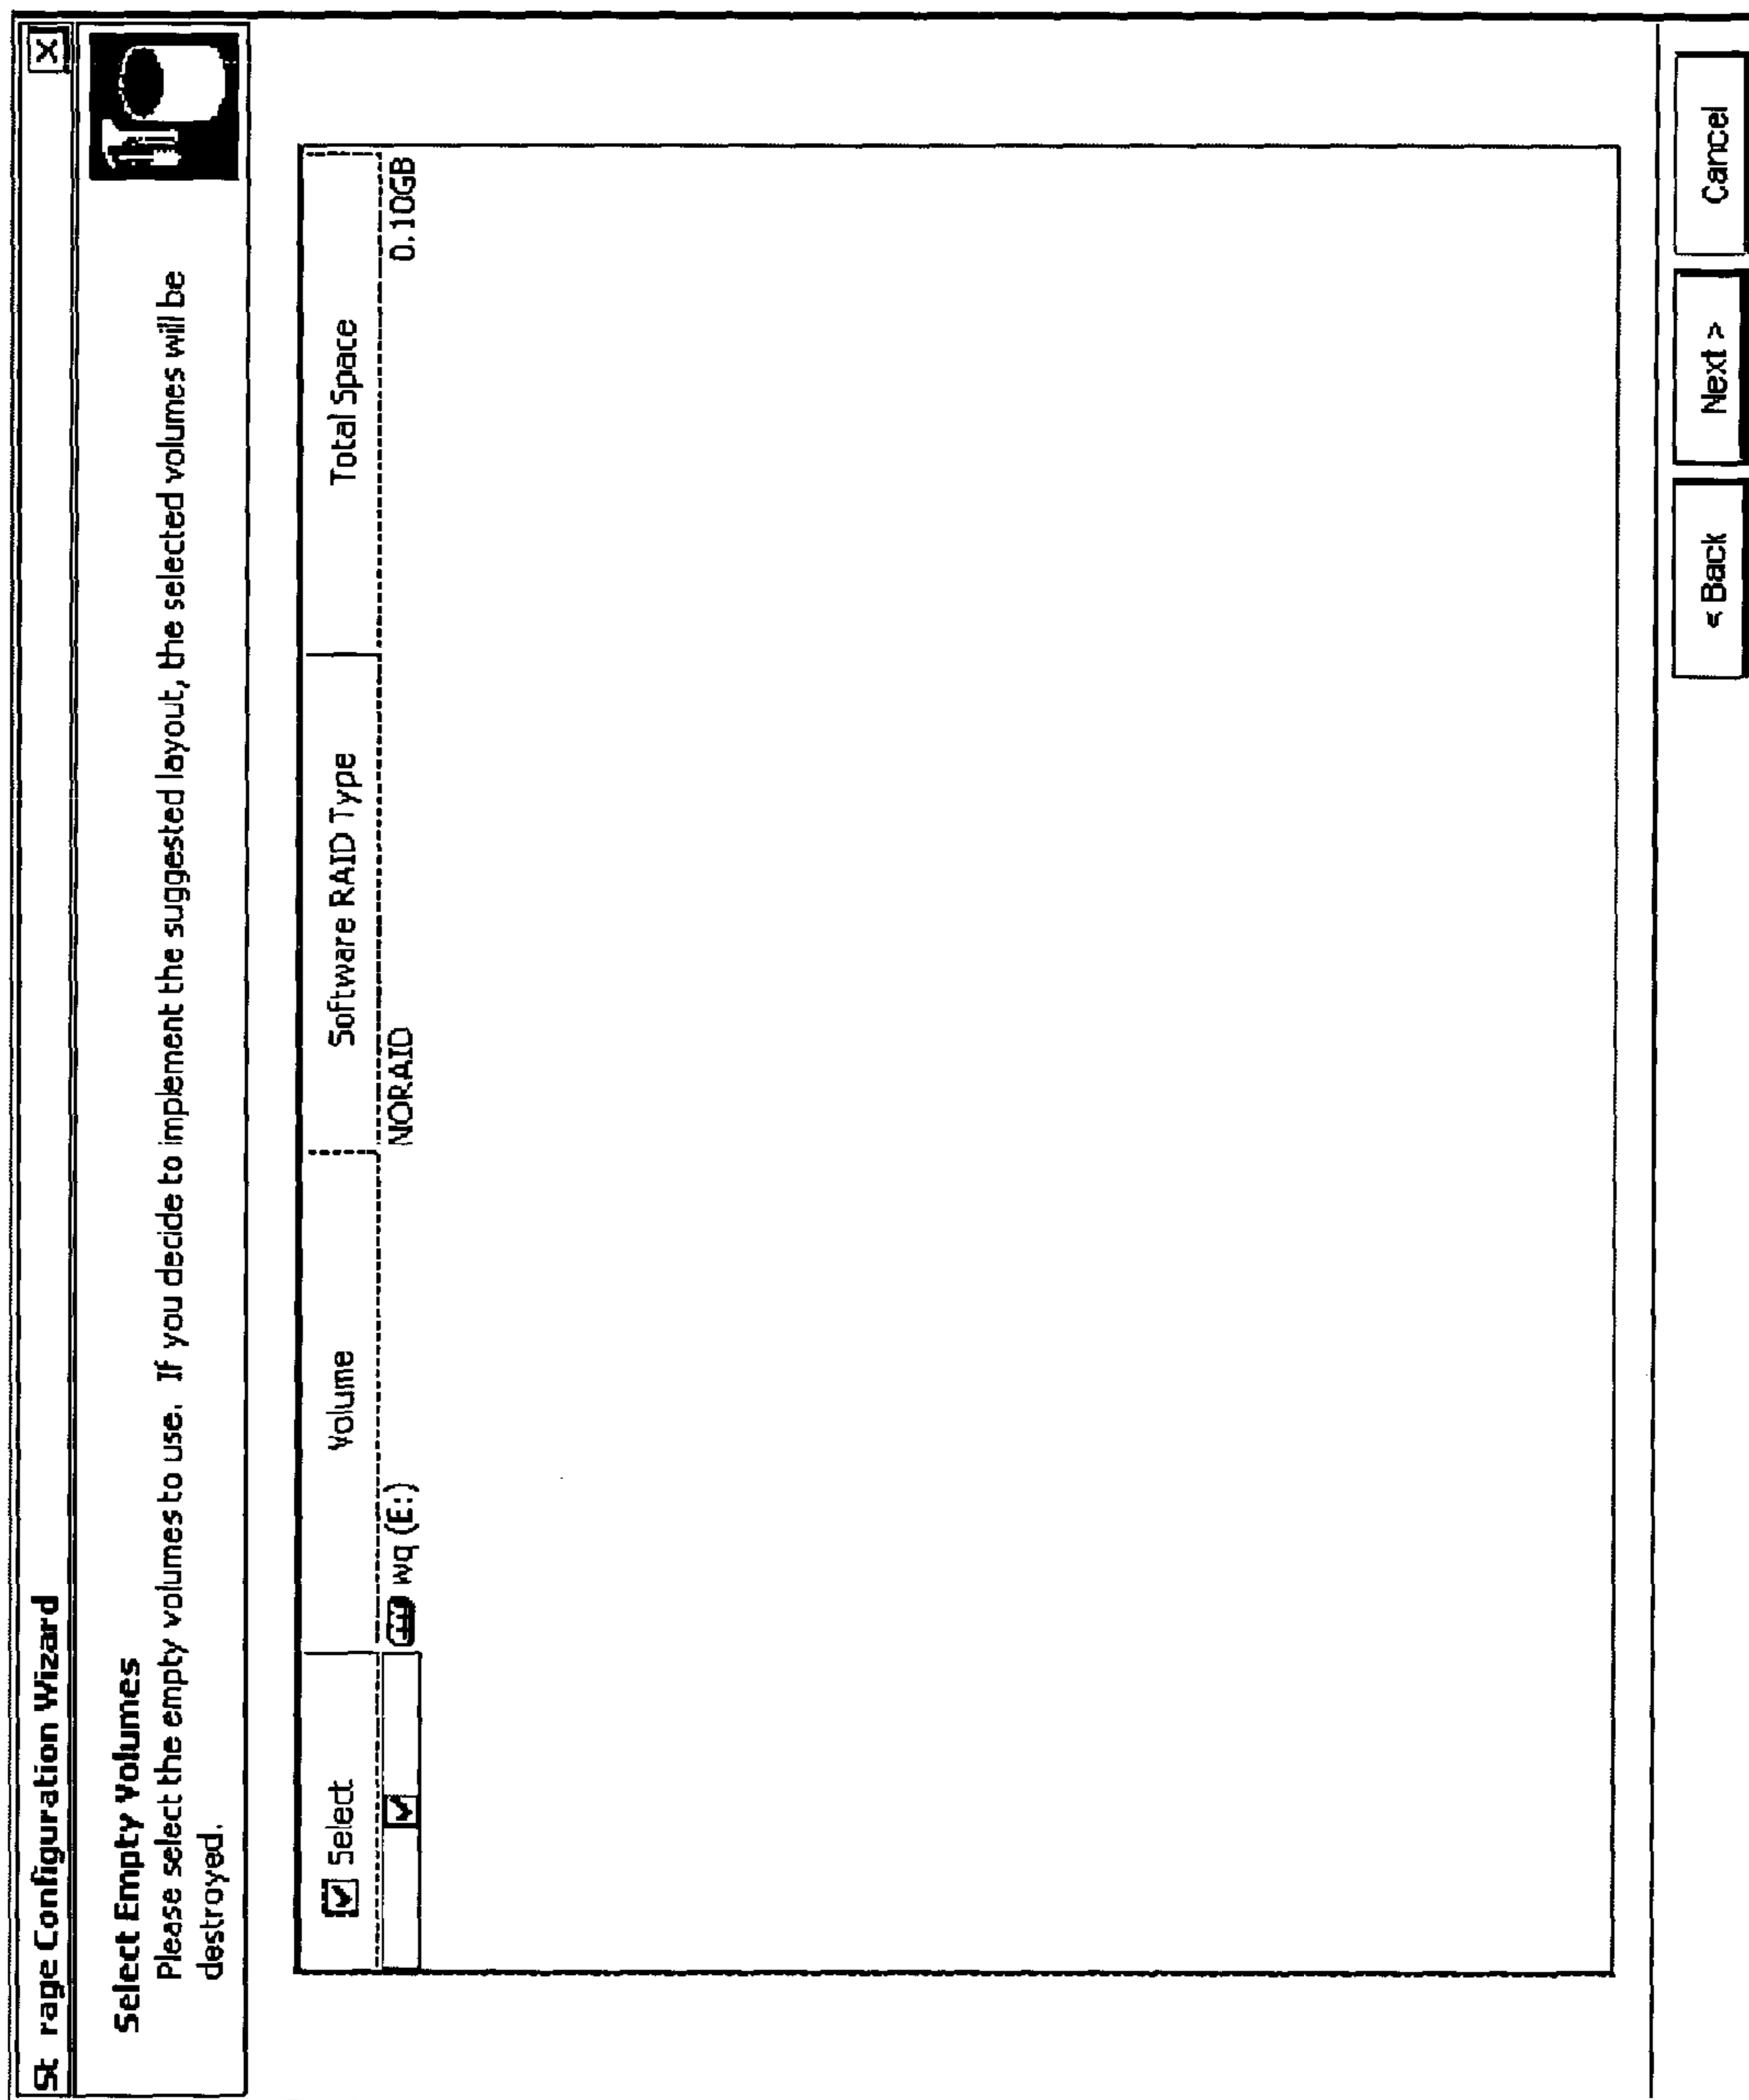


Fig. 17

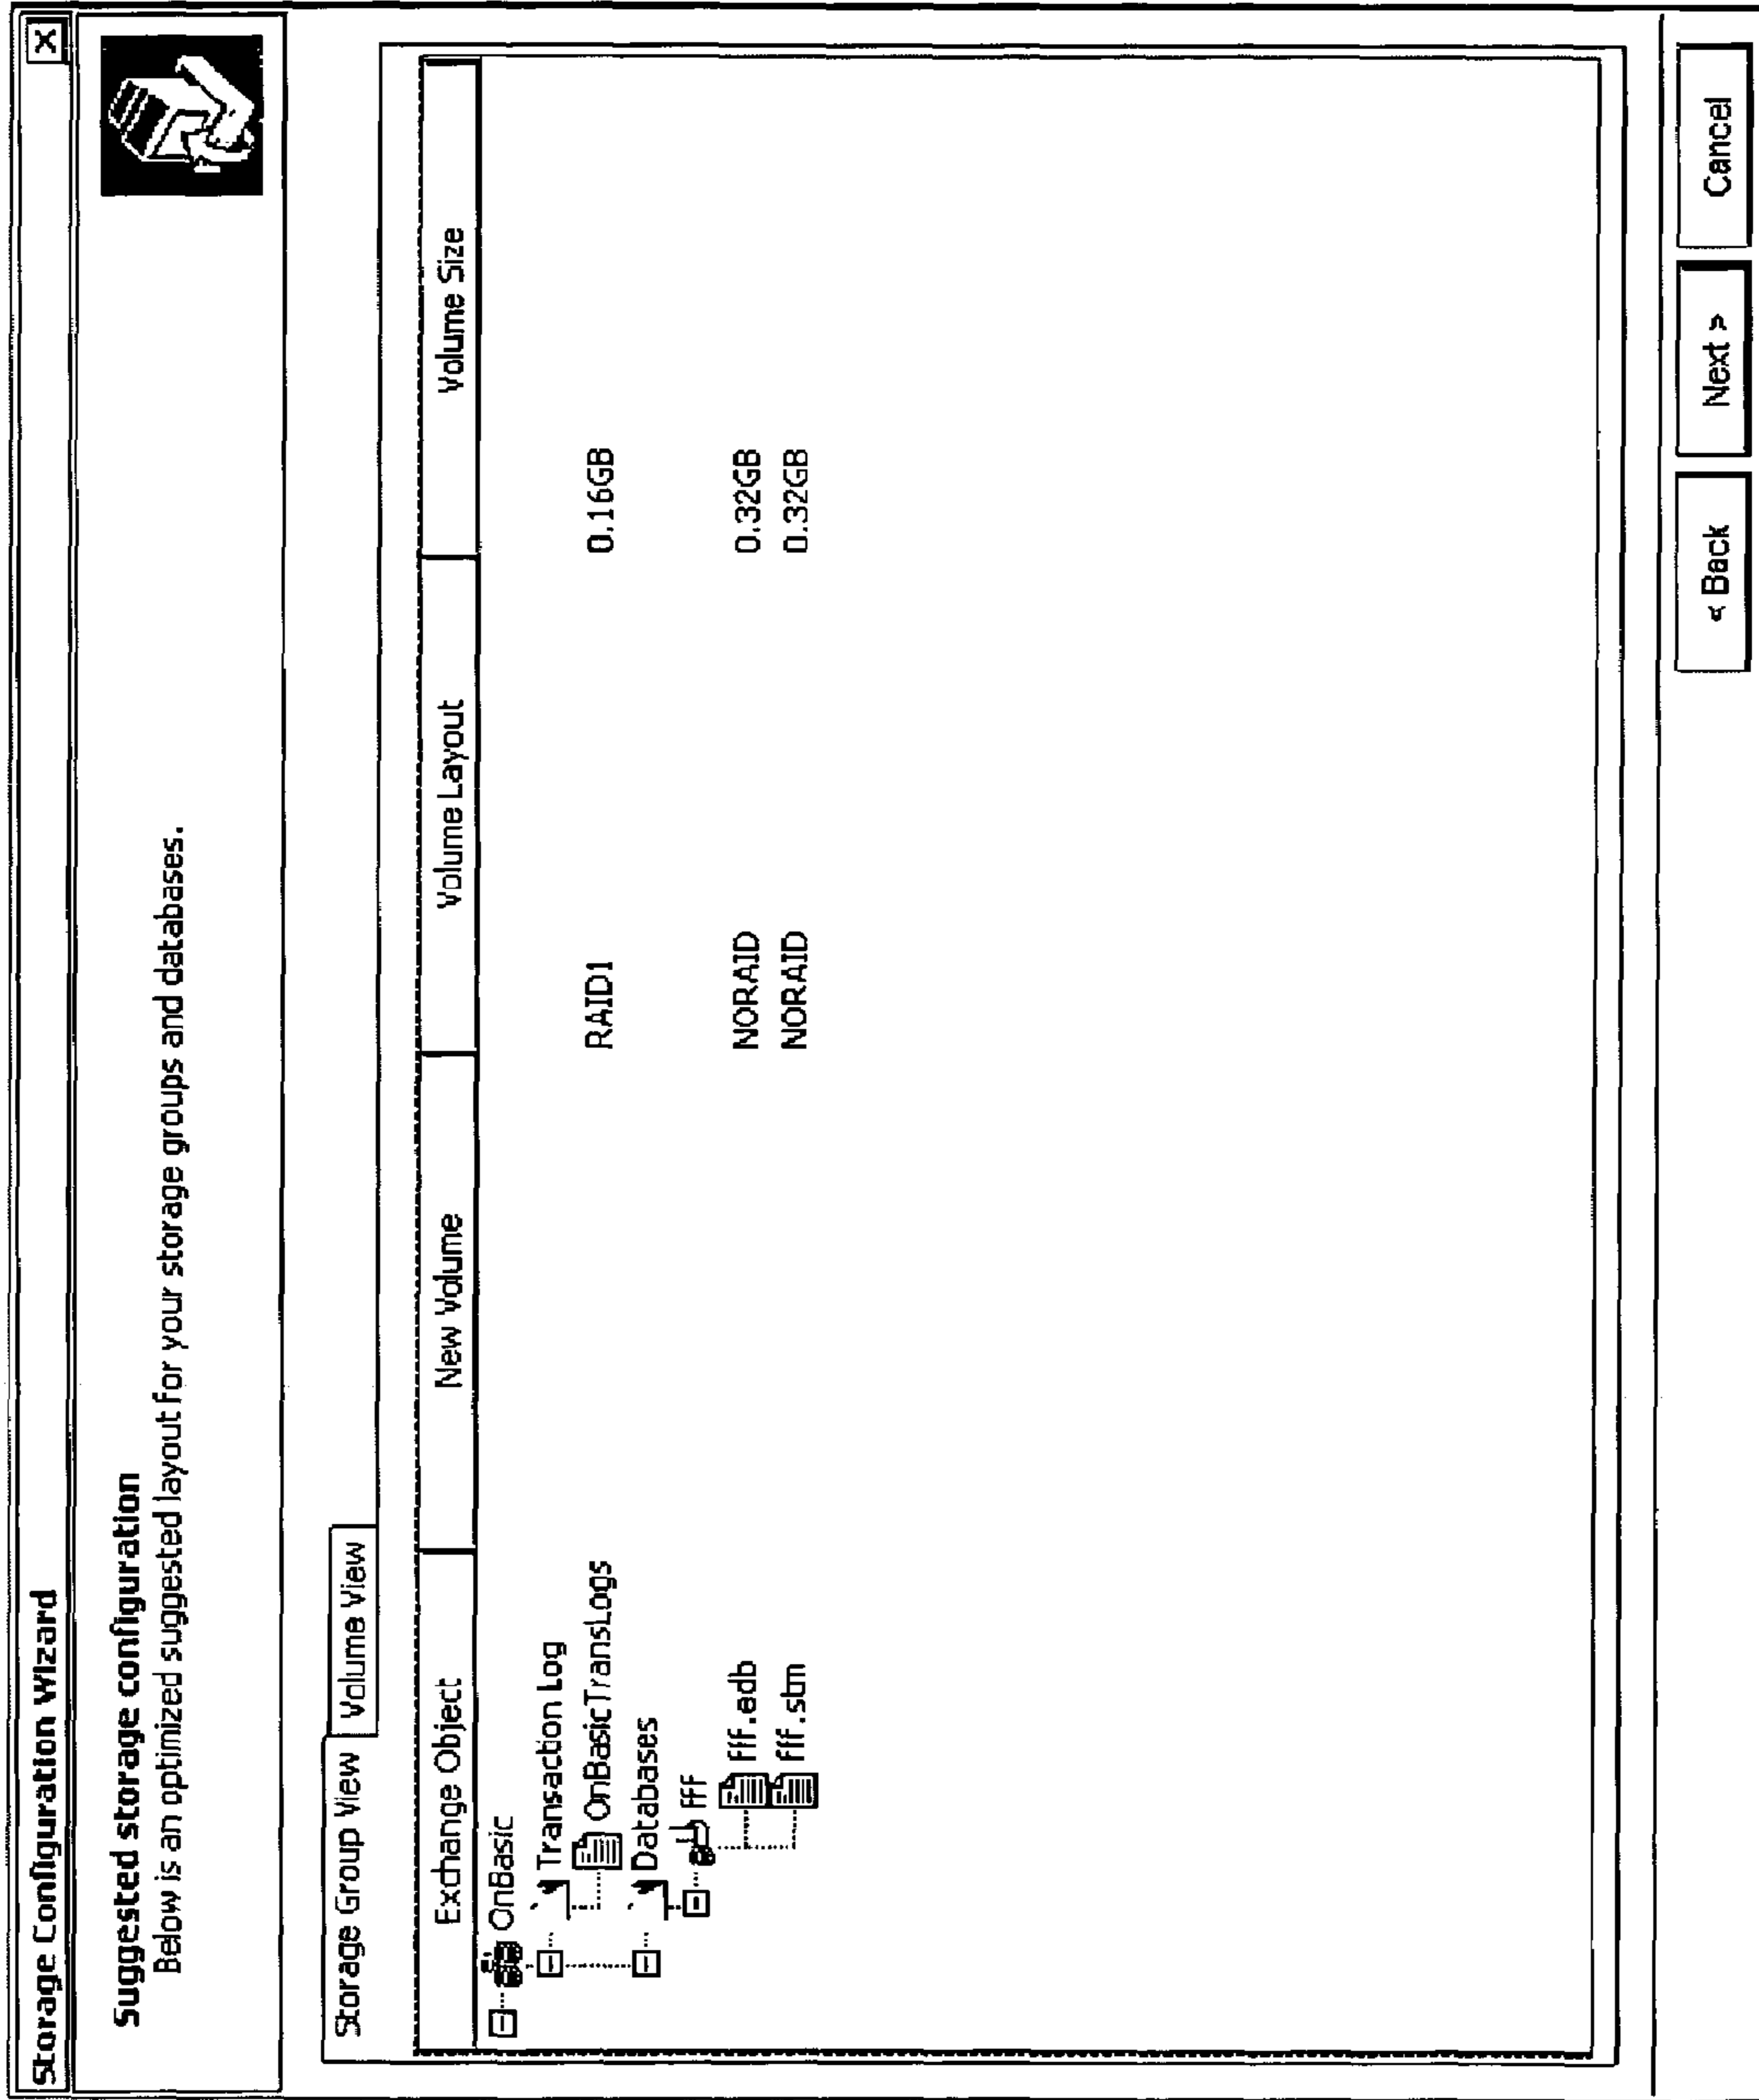


Fig. 18

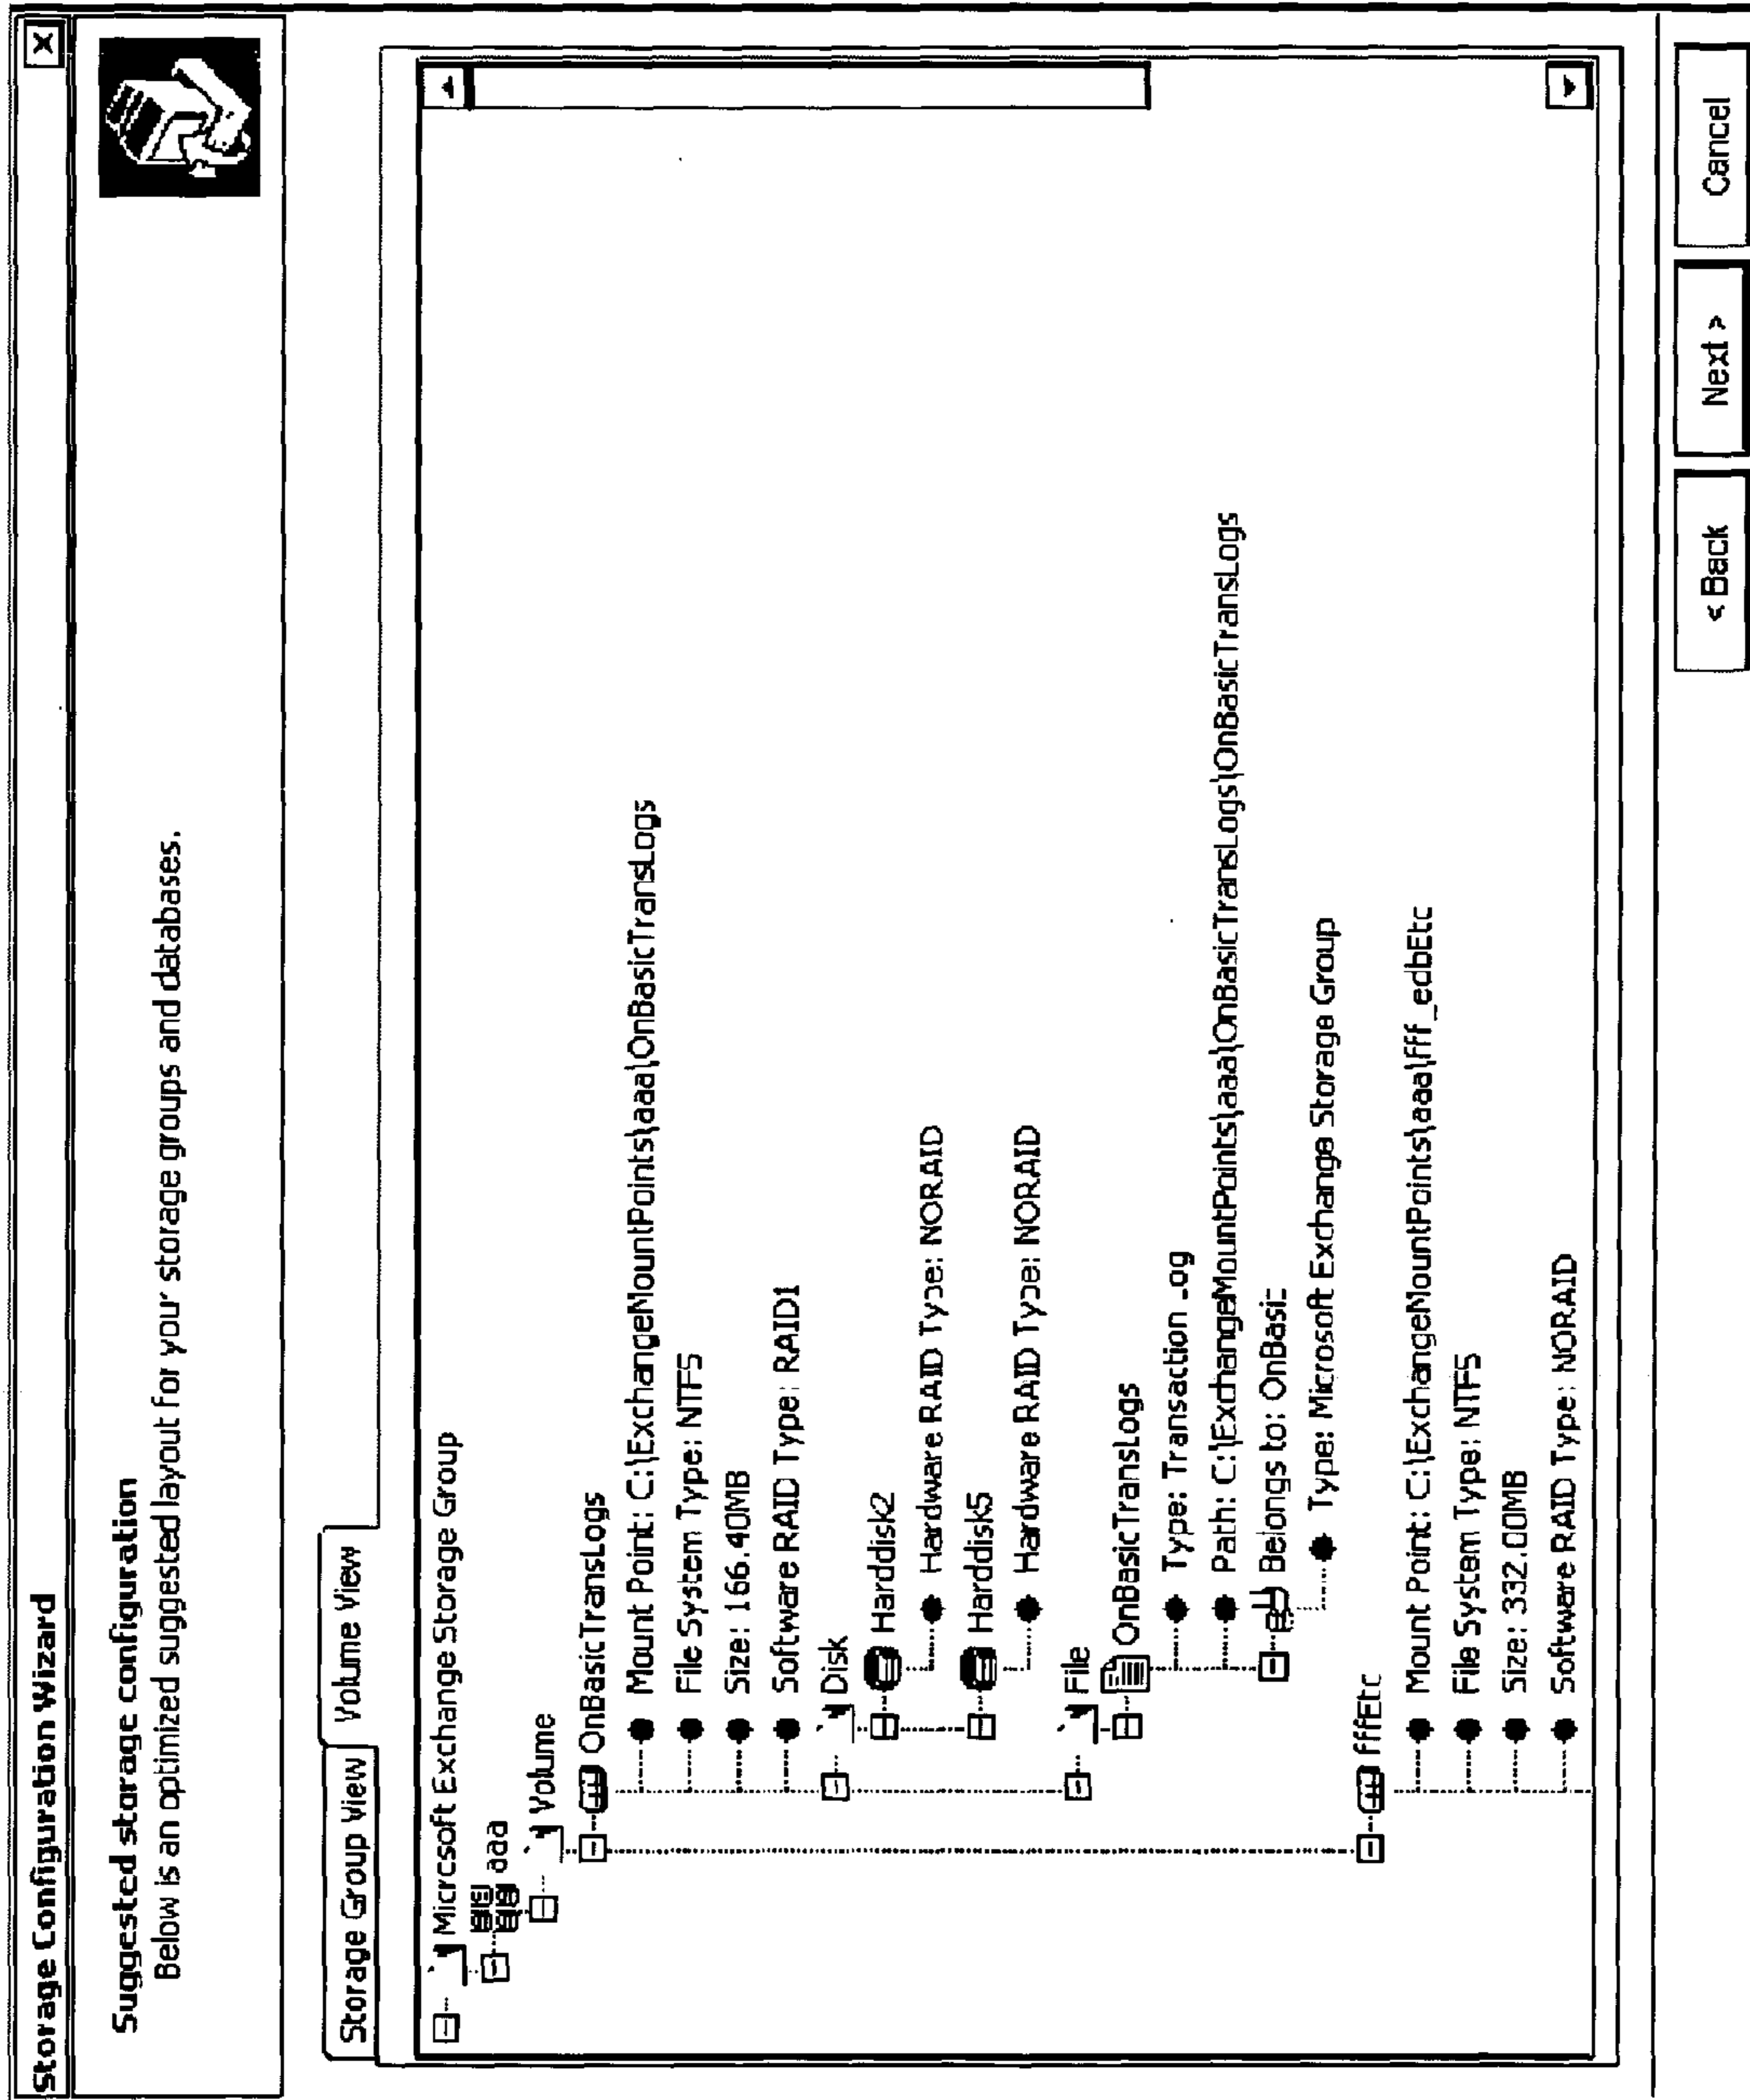


Fig. 19



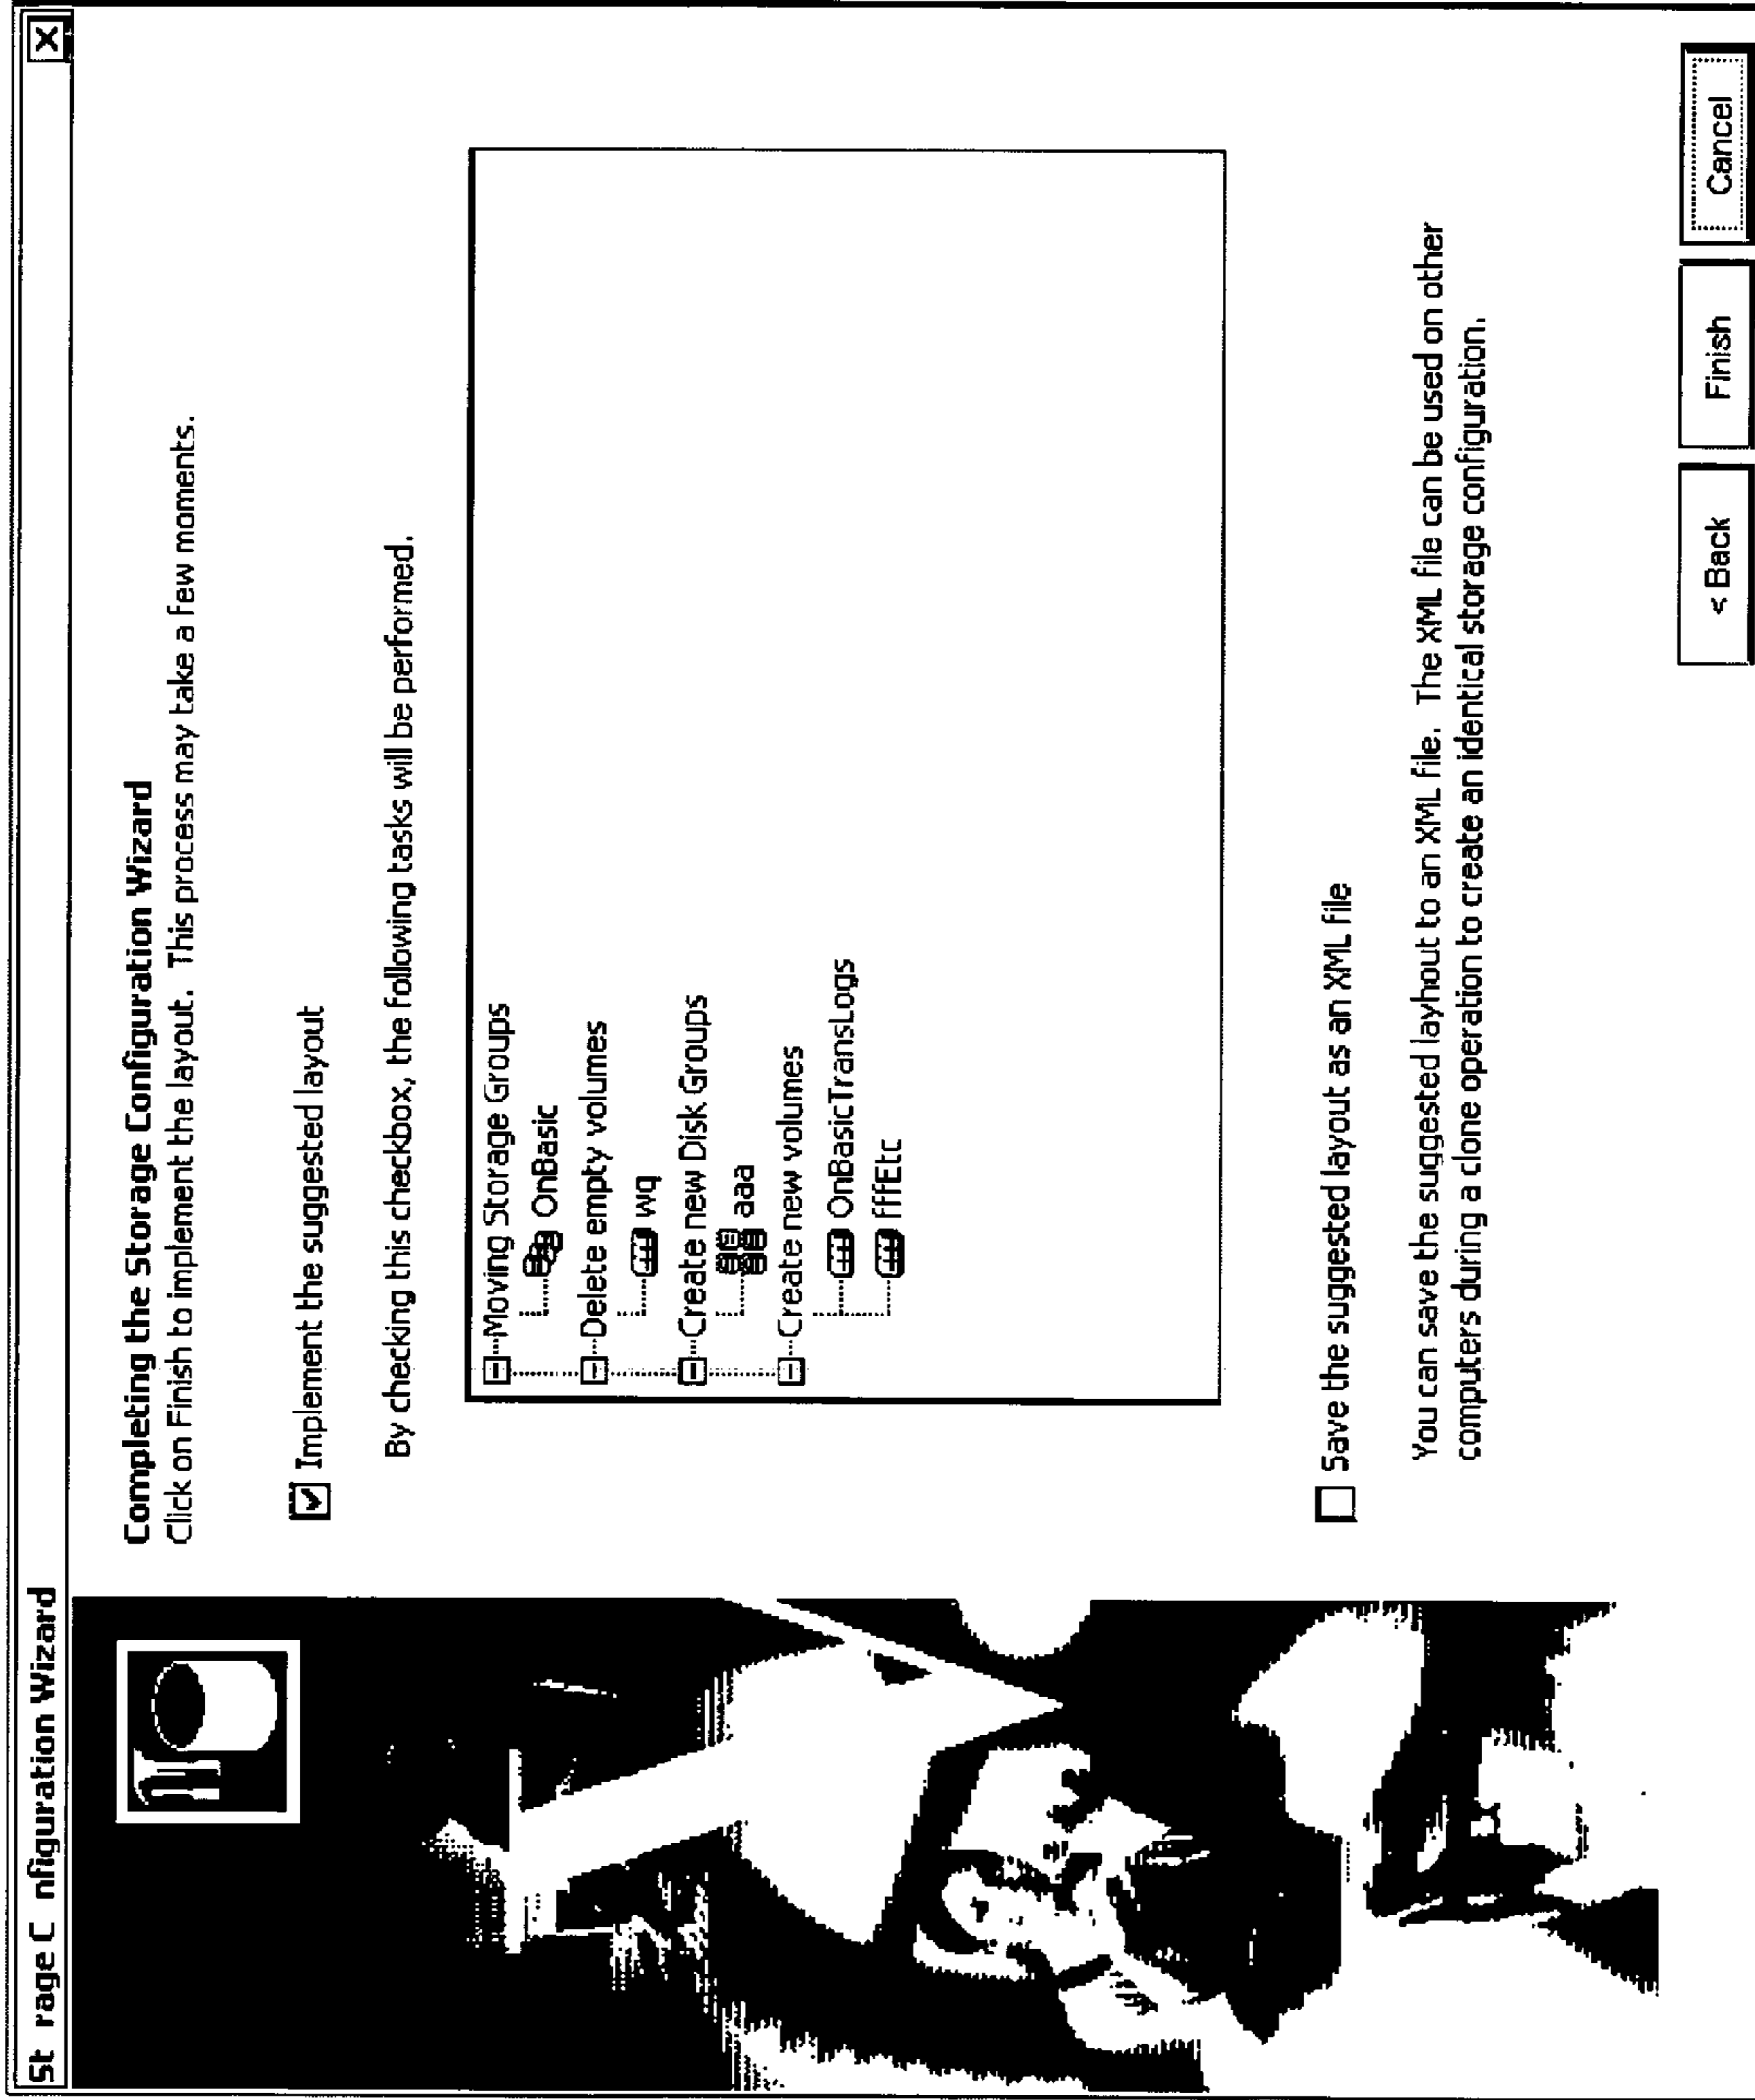


Fig. 20

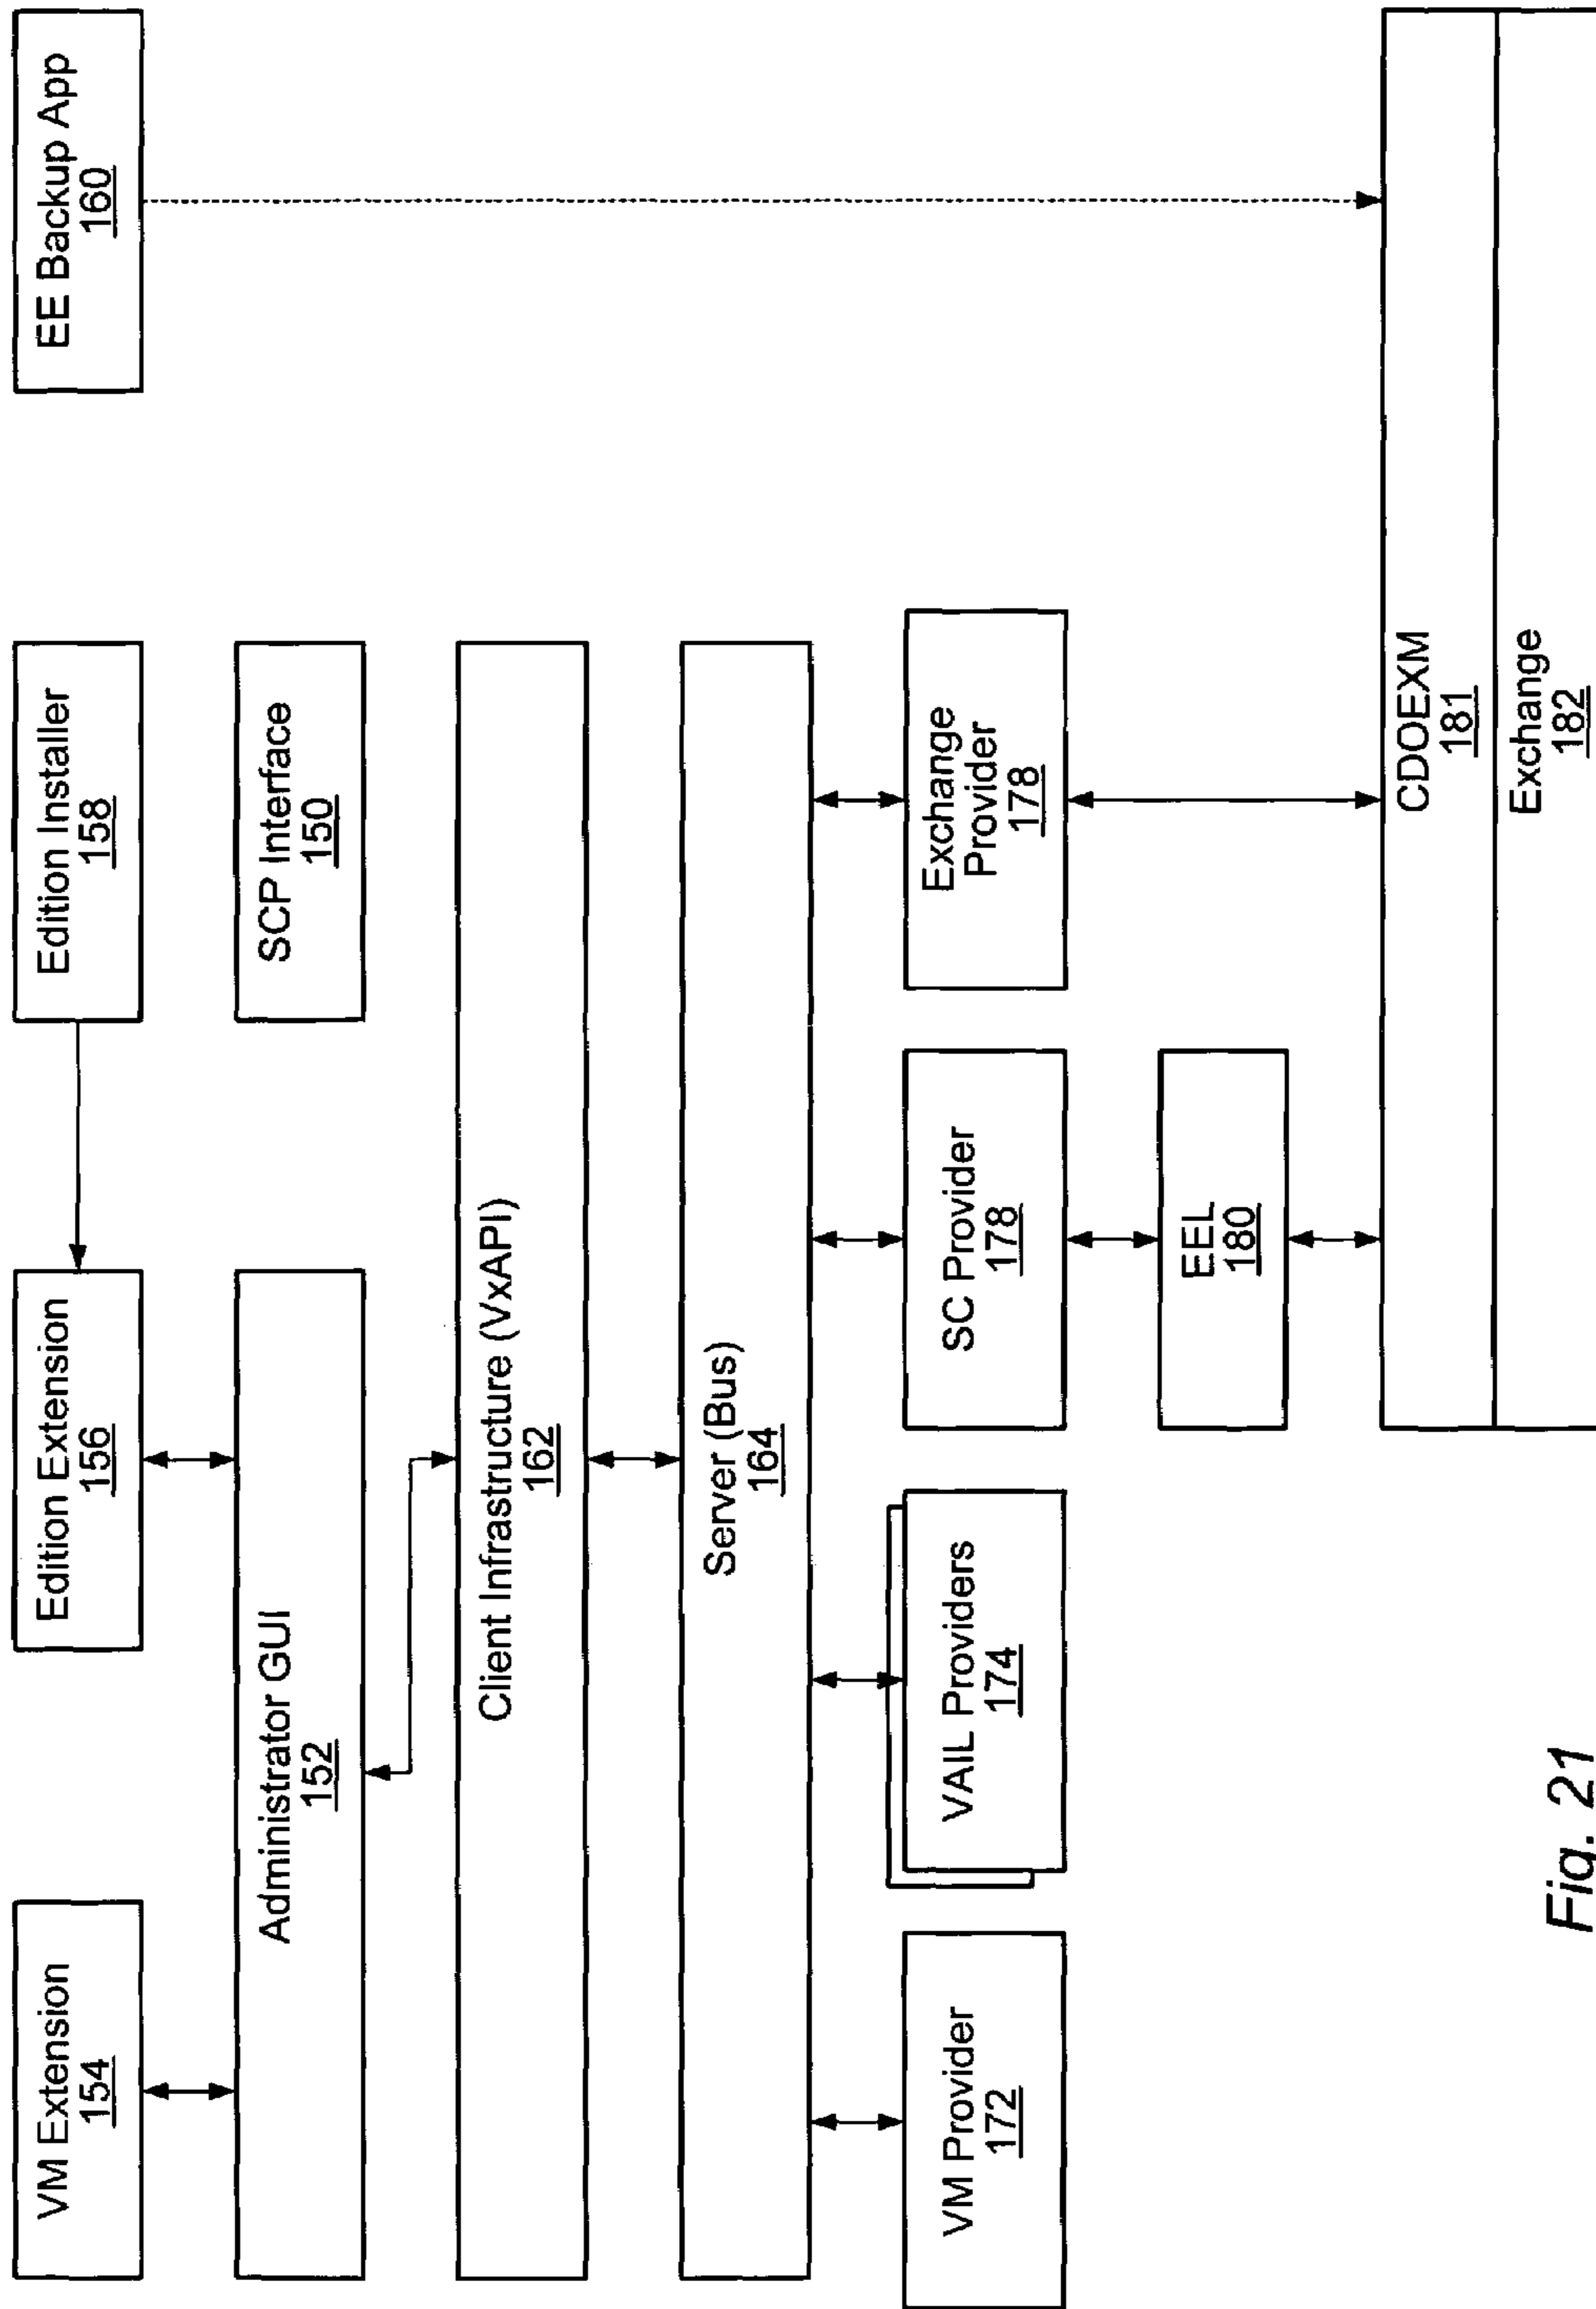


Fig. 21

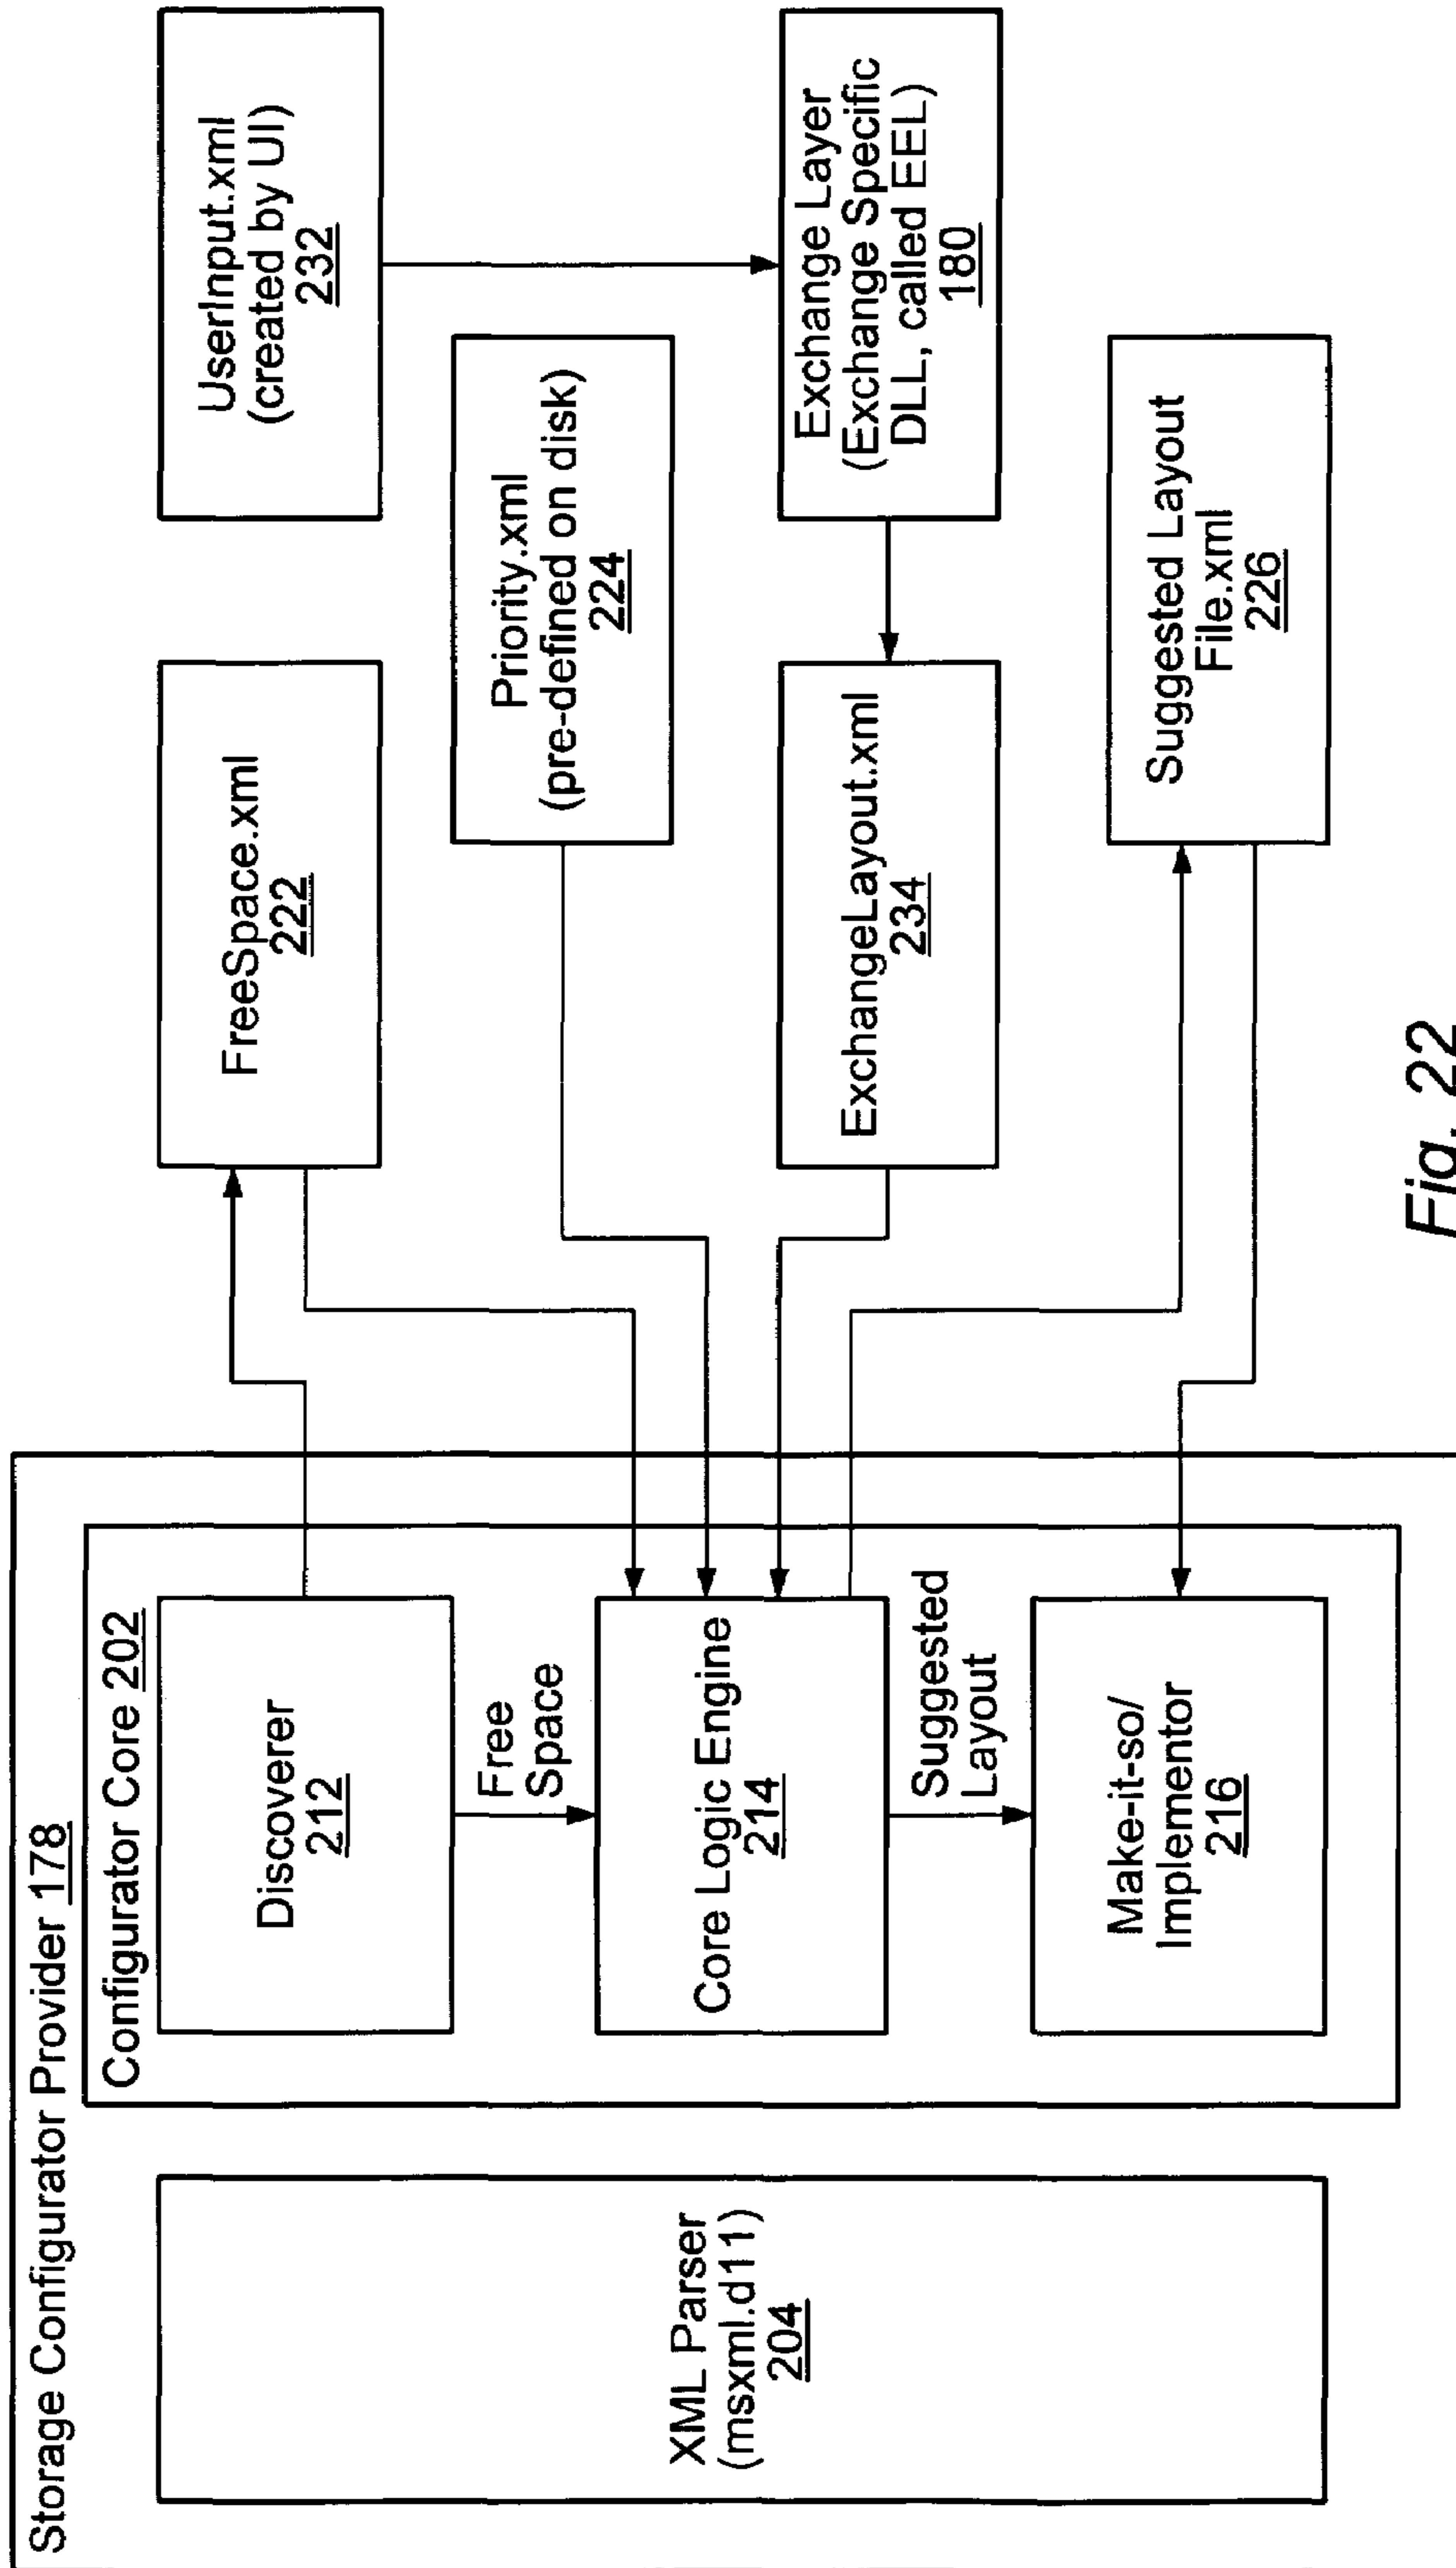


Fig. 22

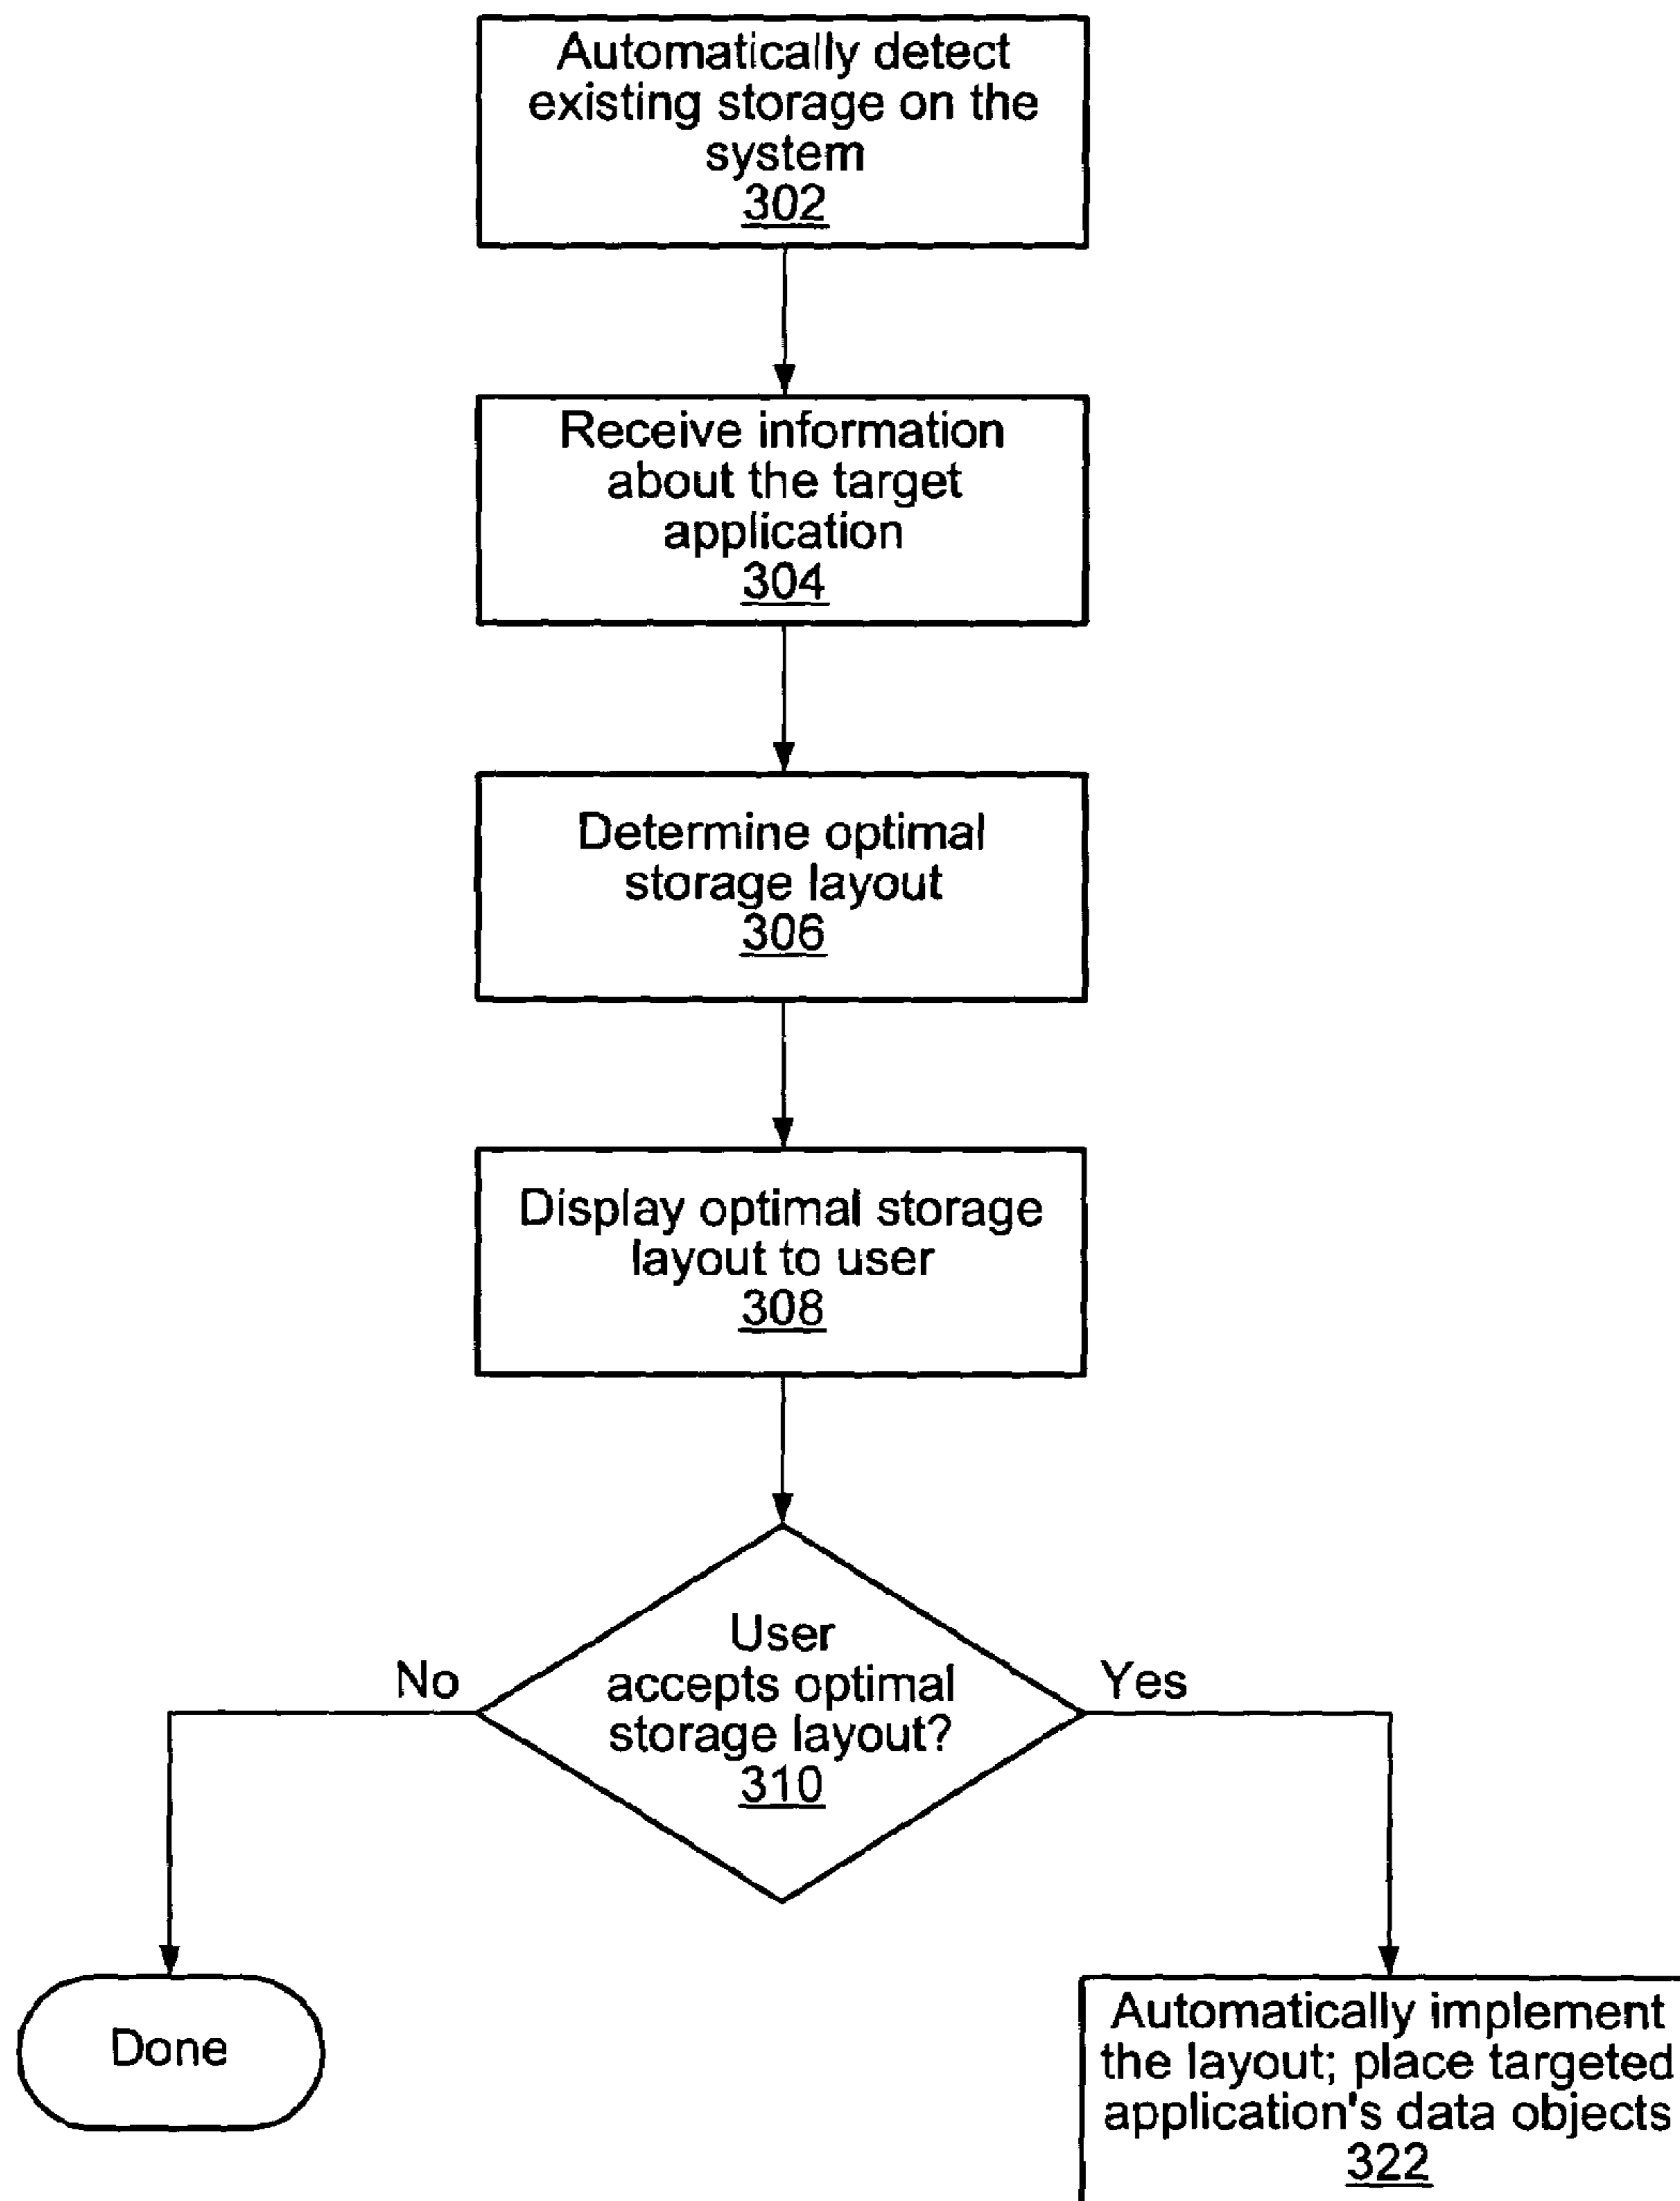


Fig. 23

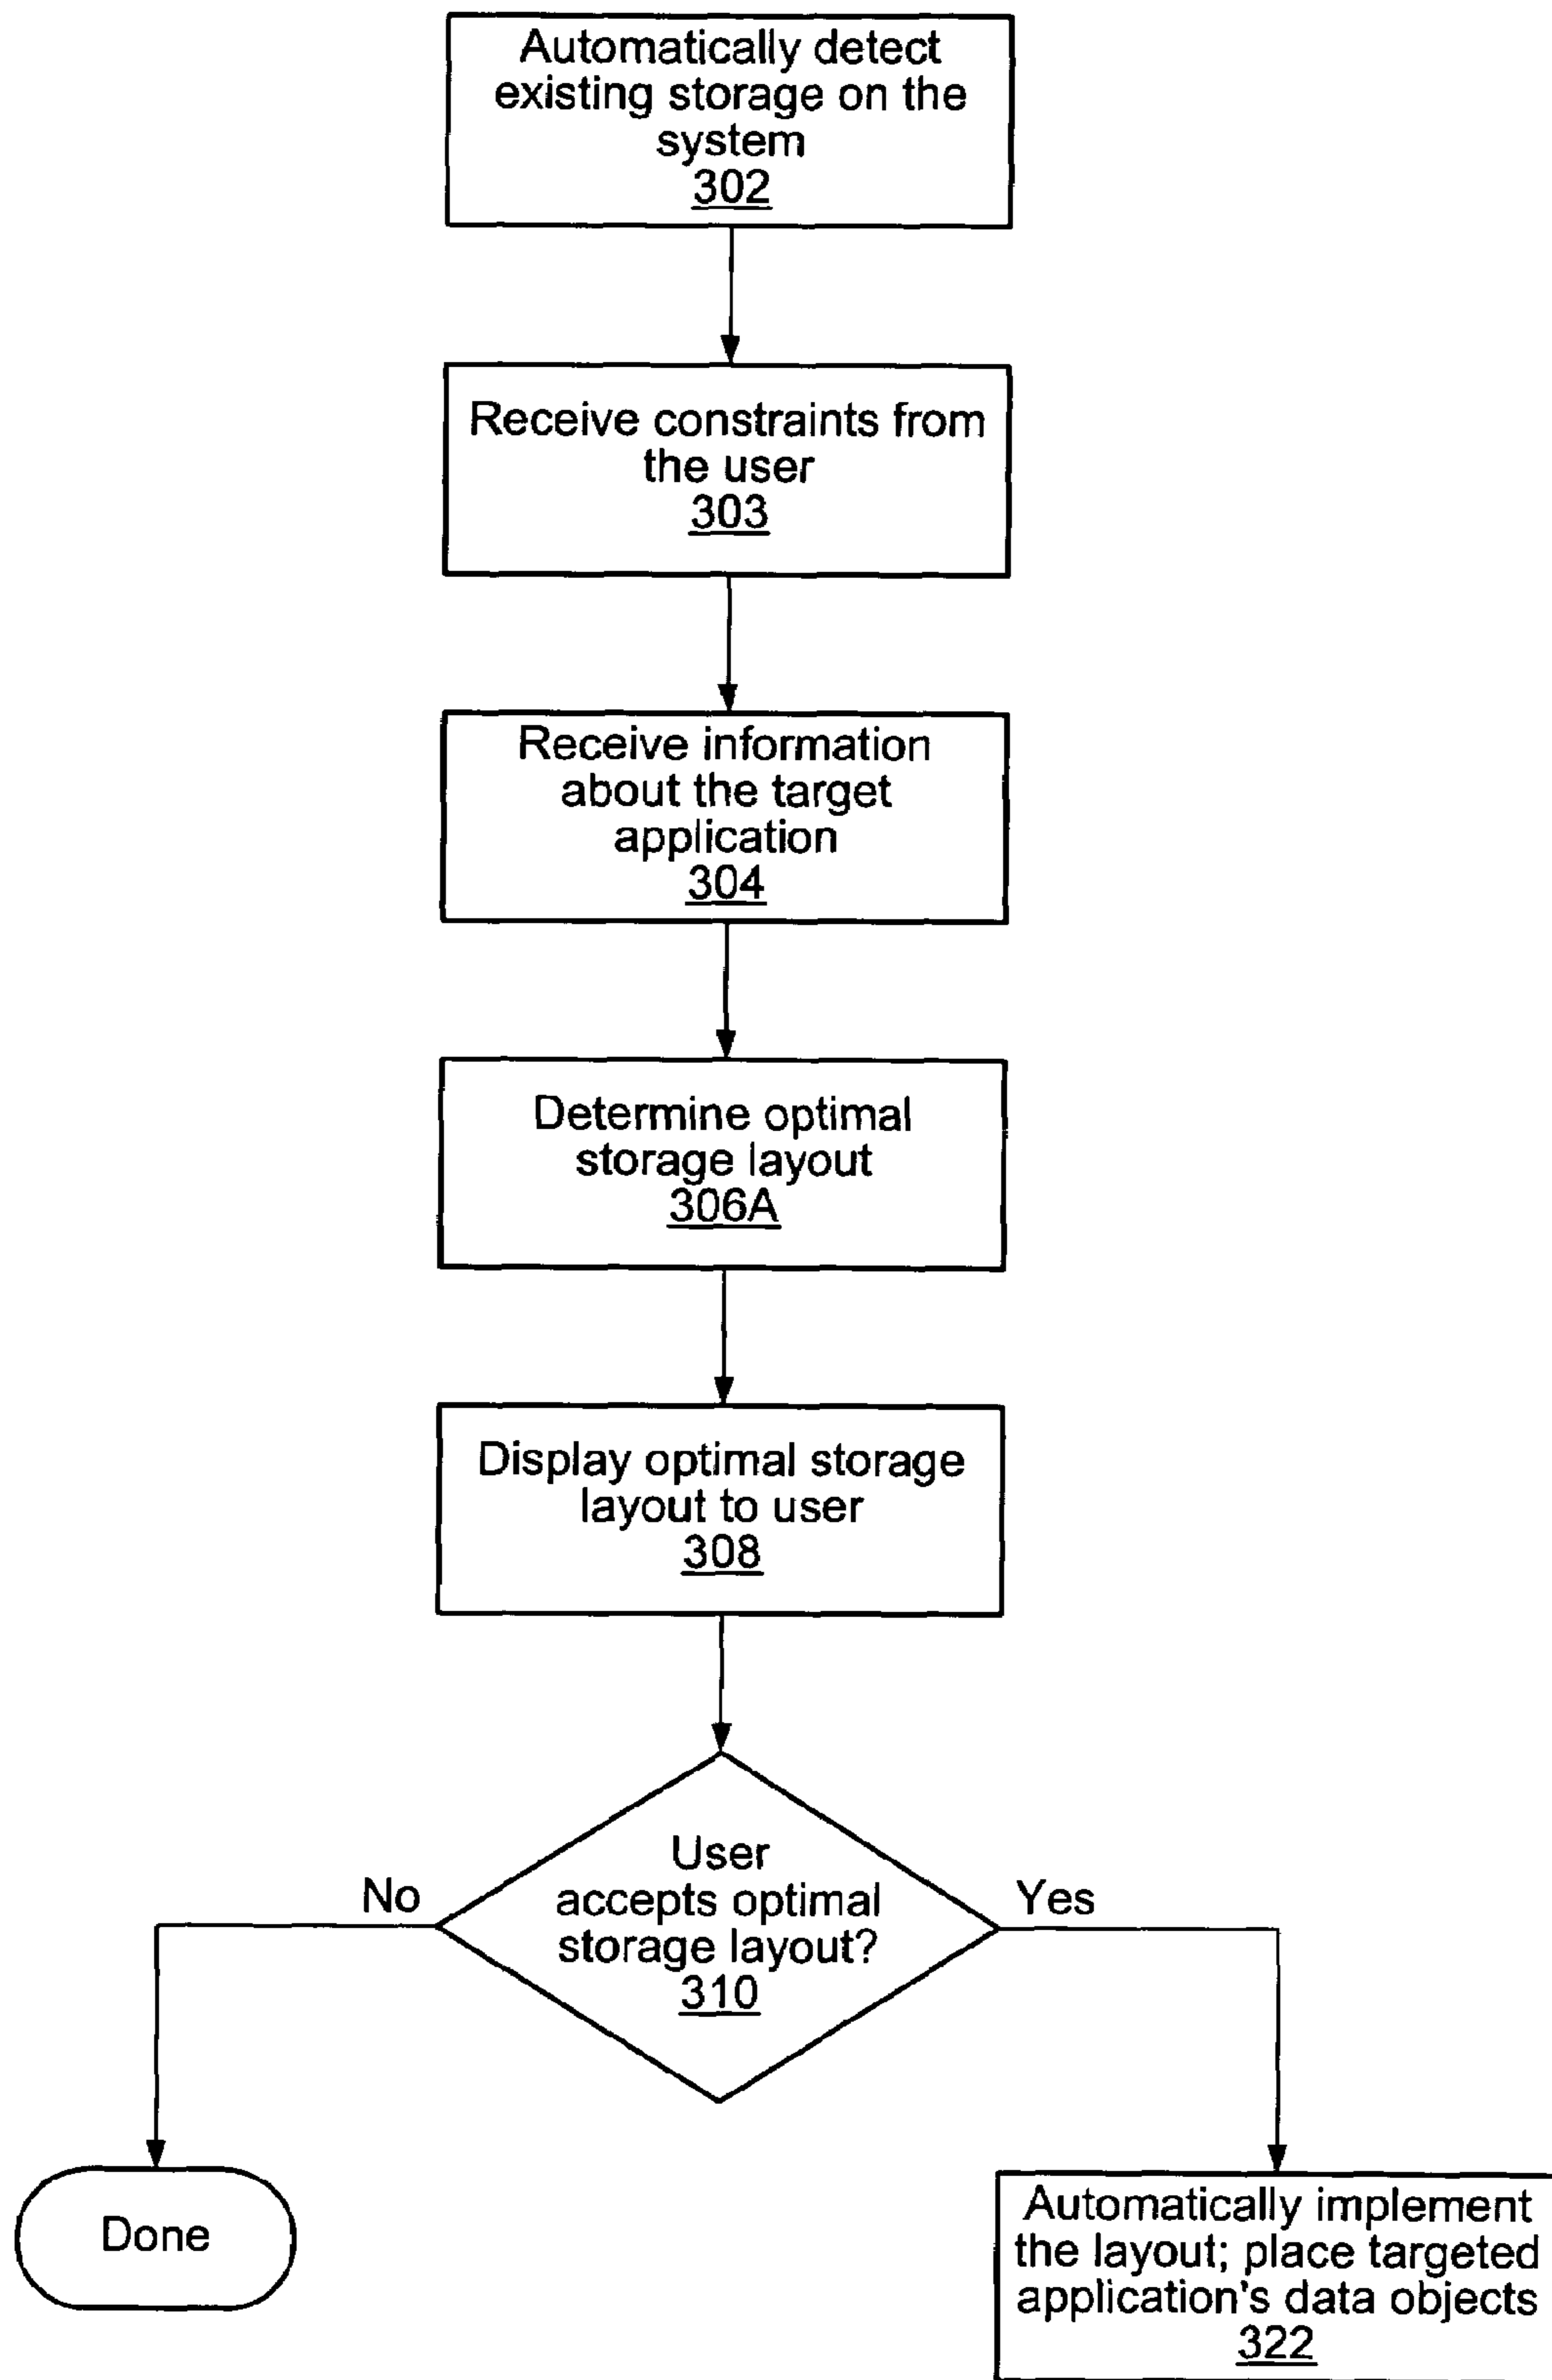
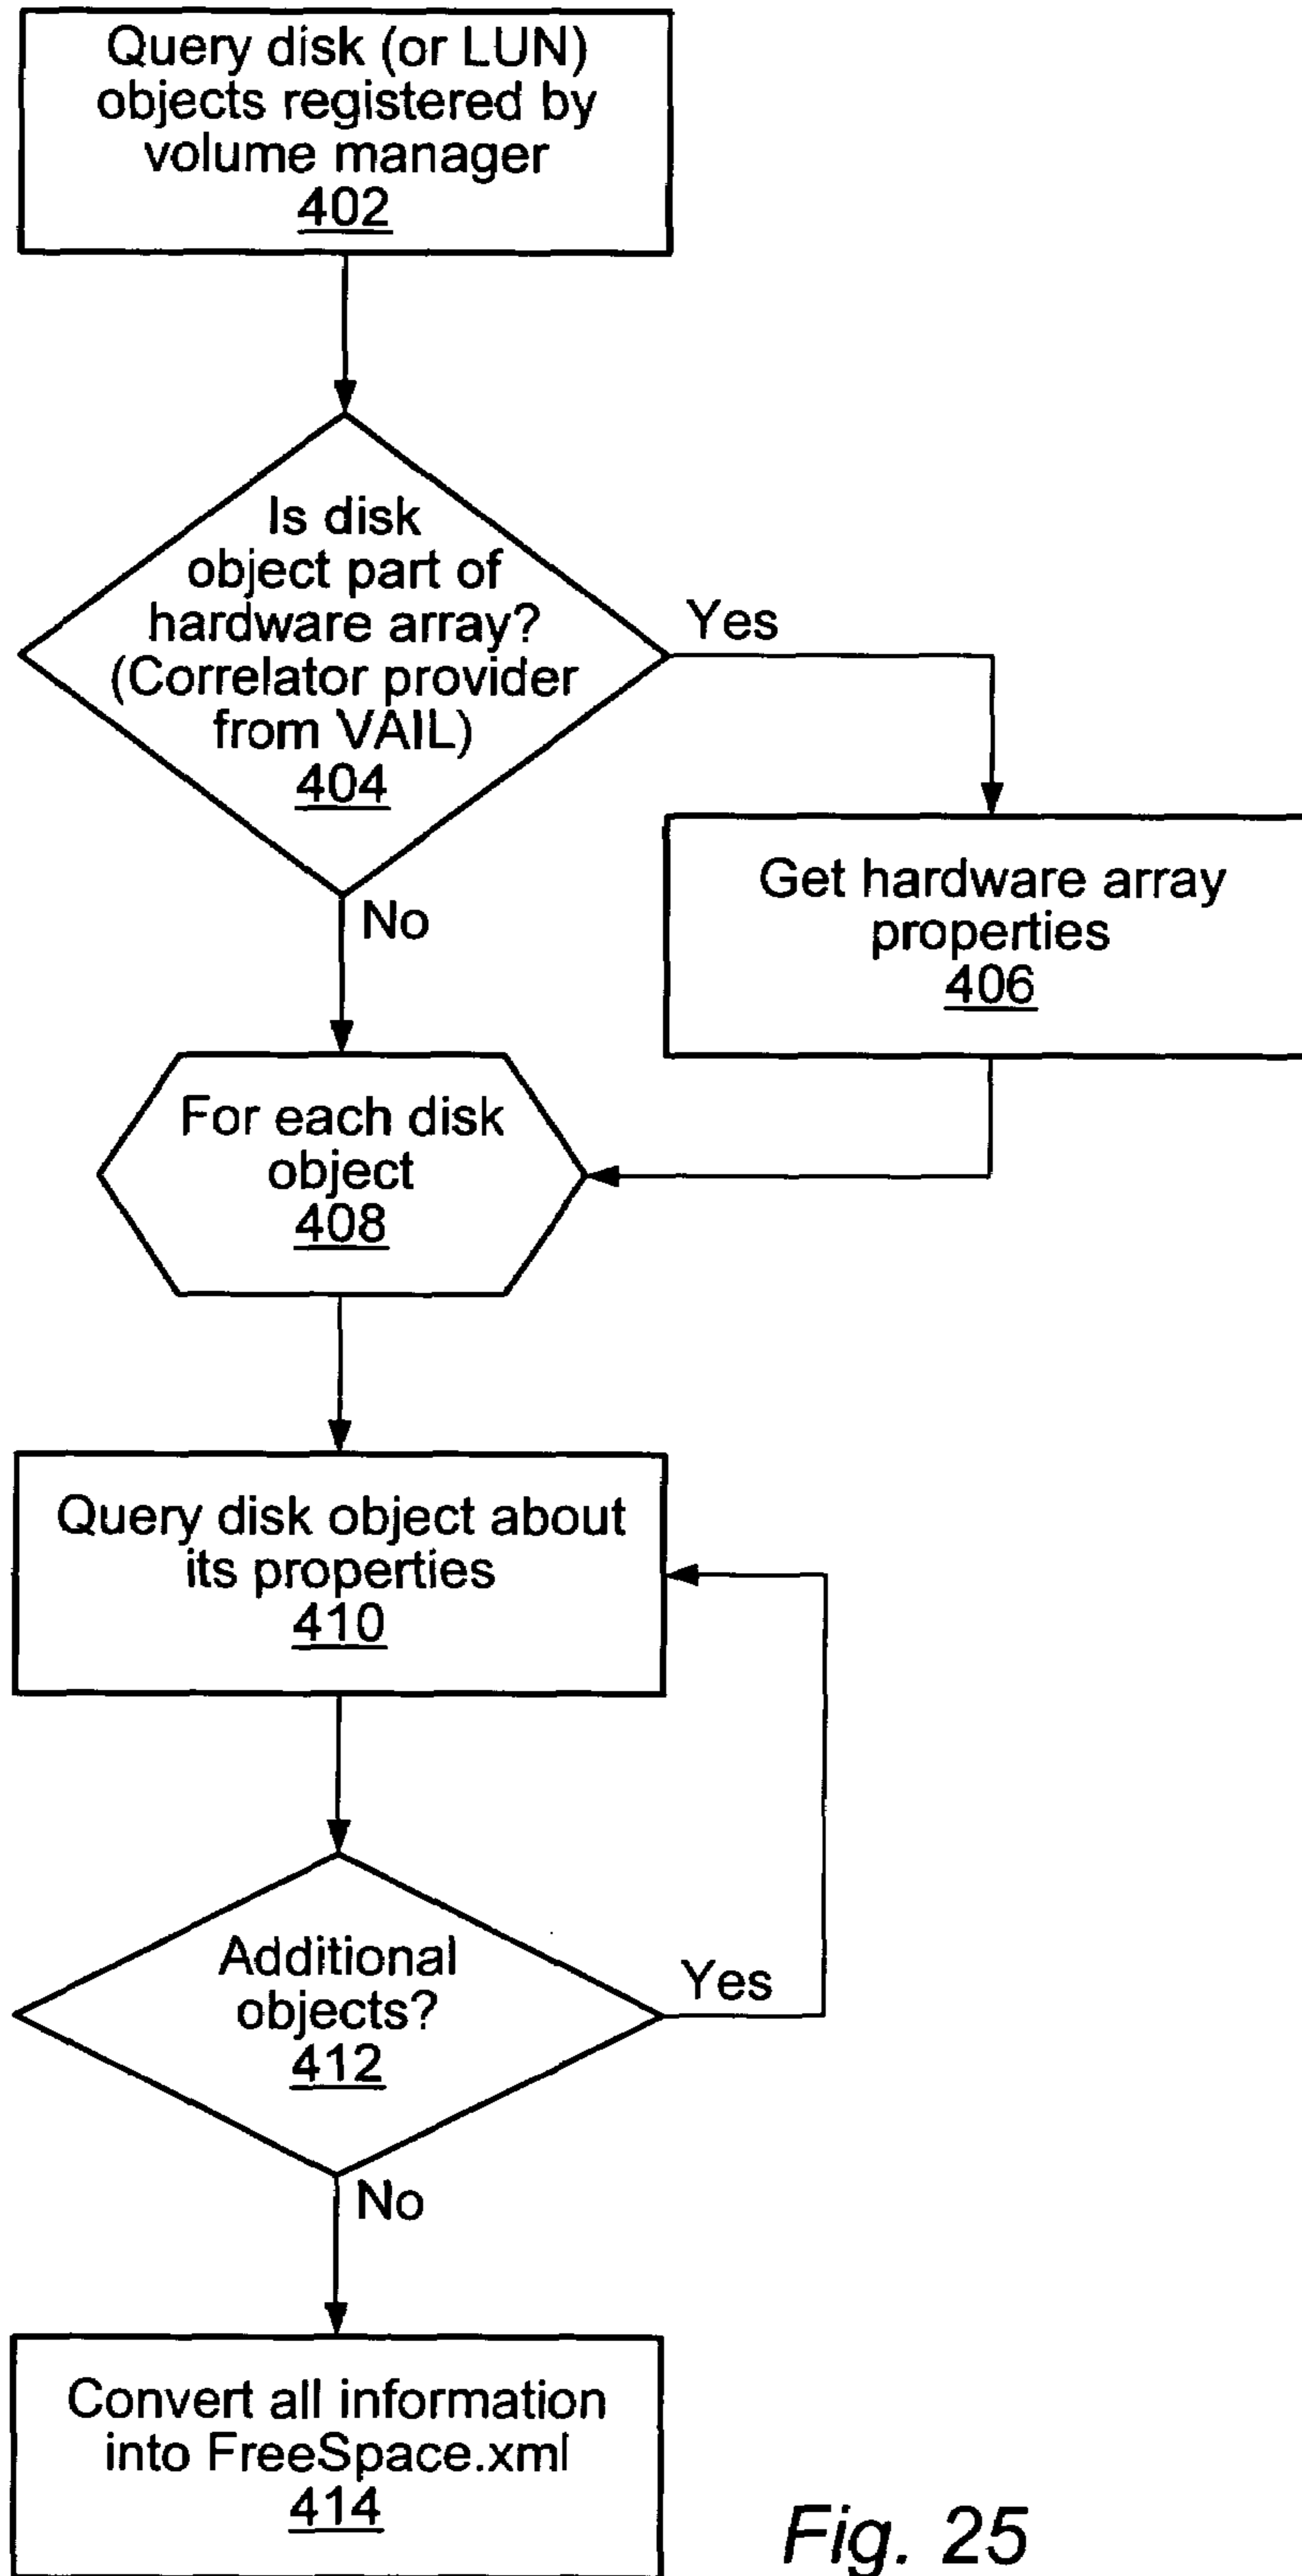


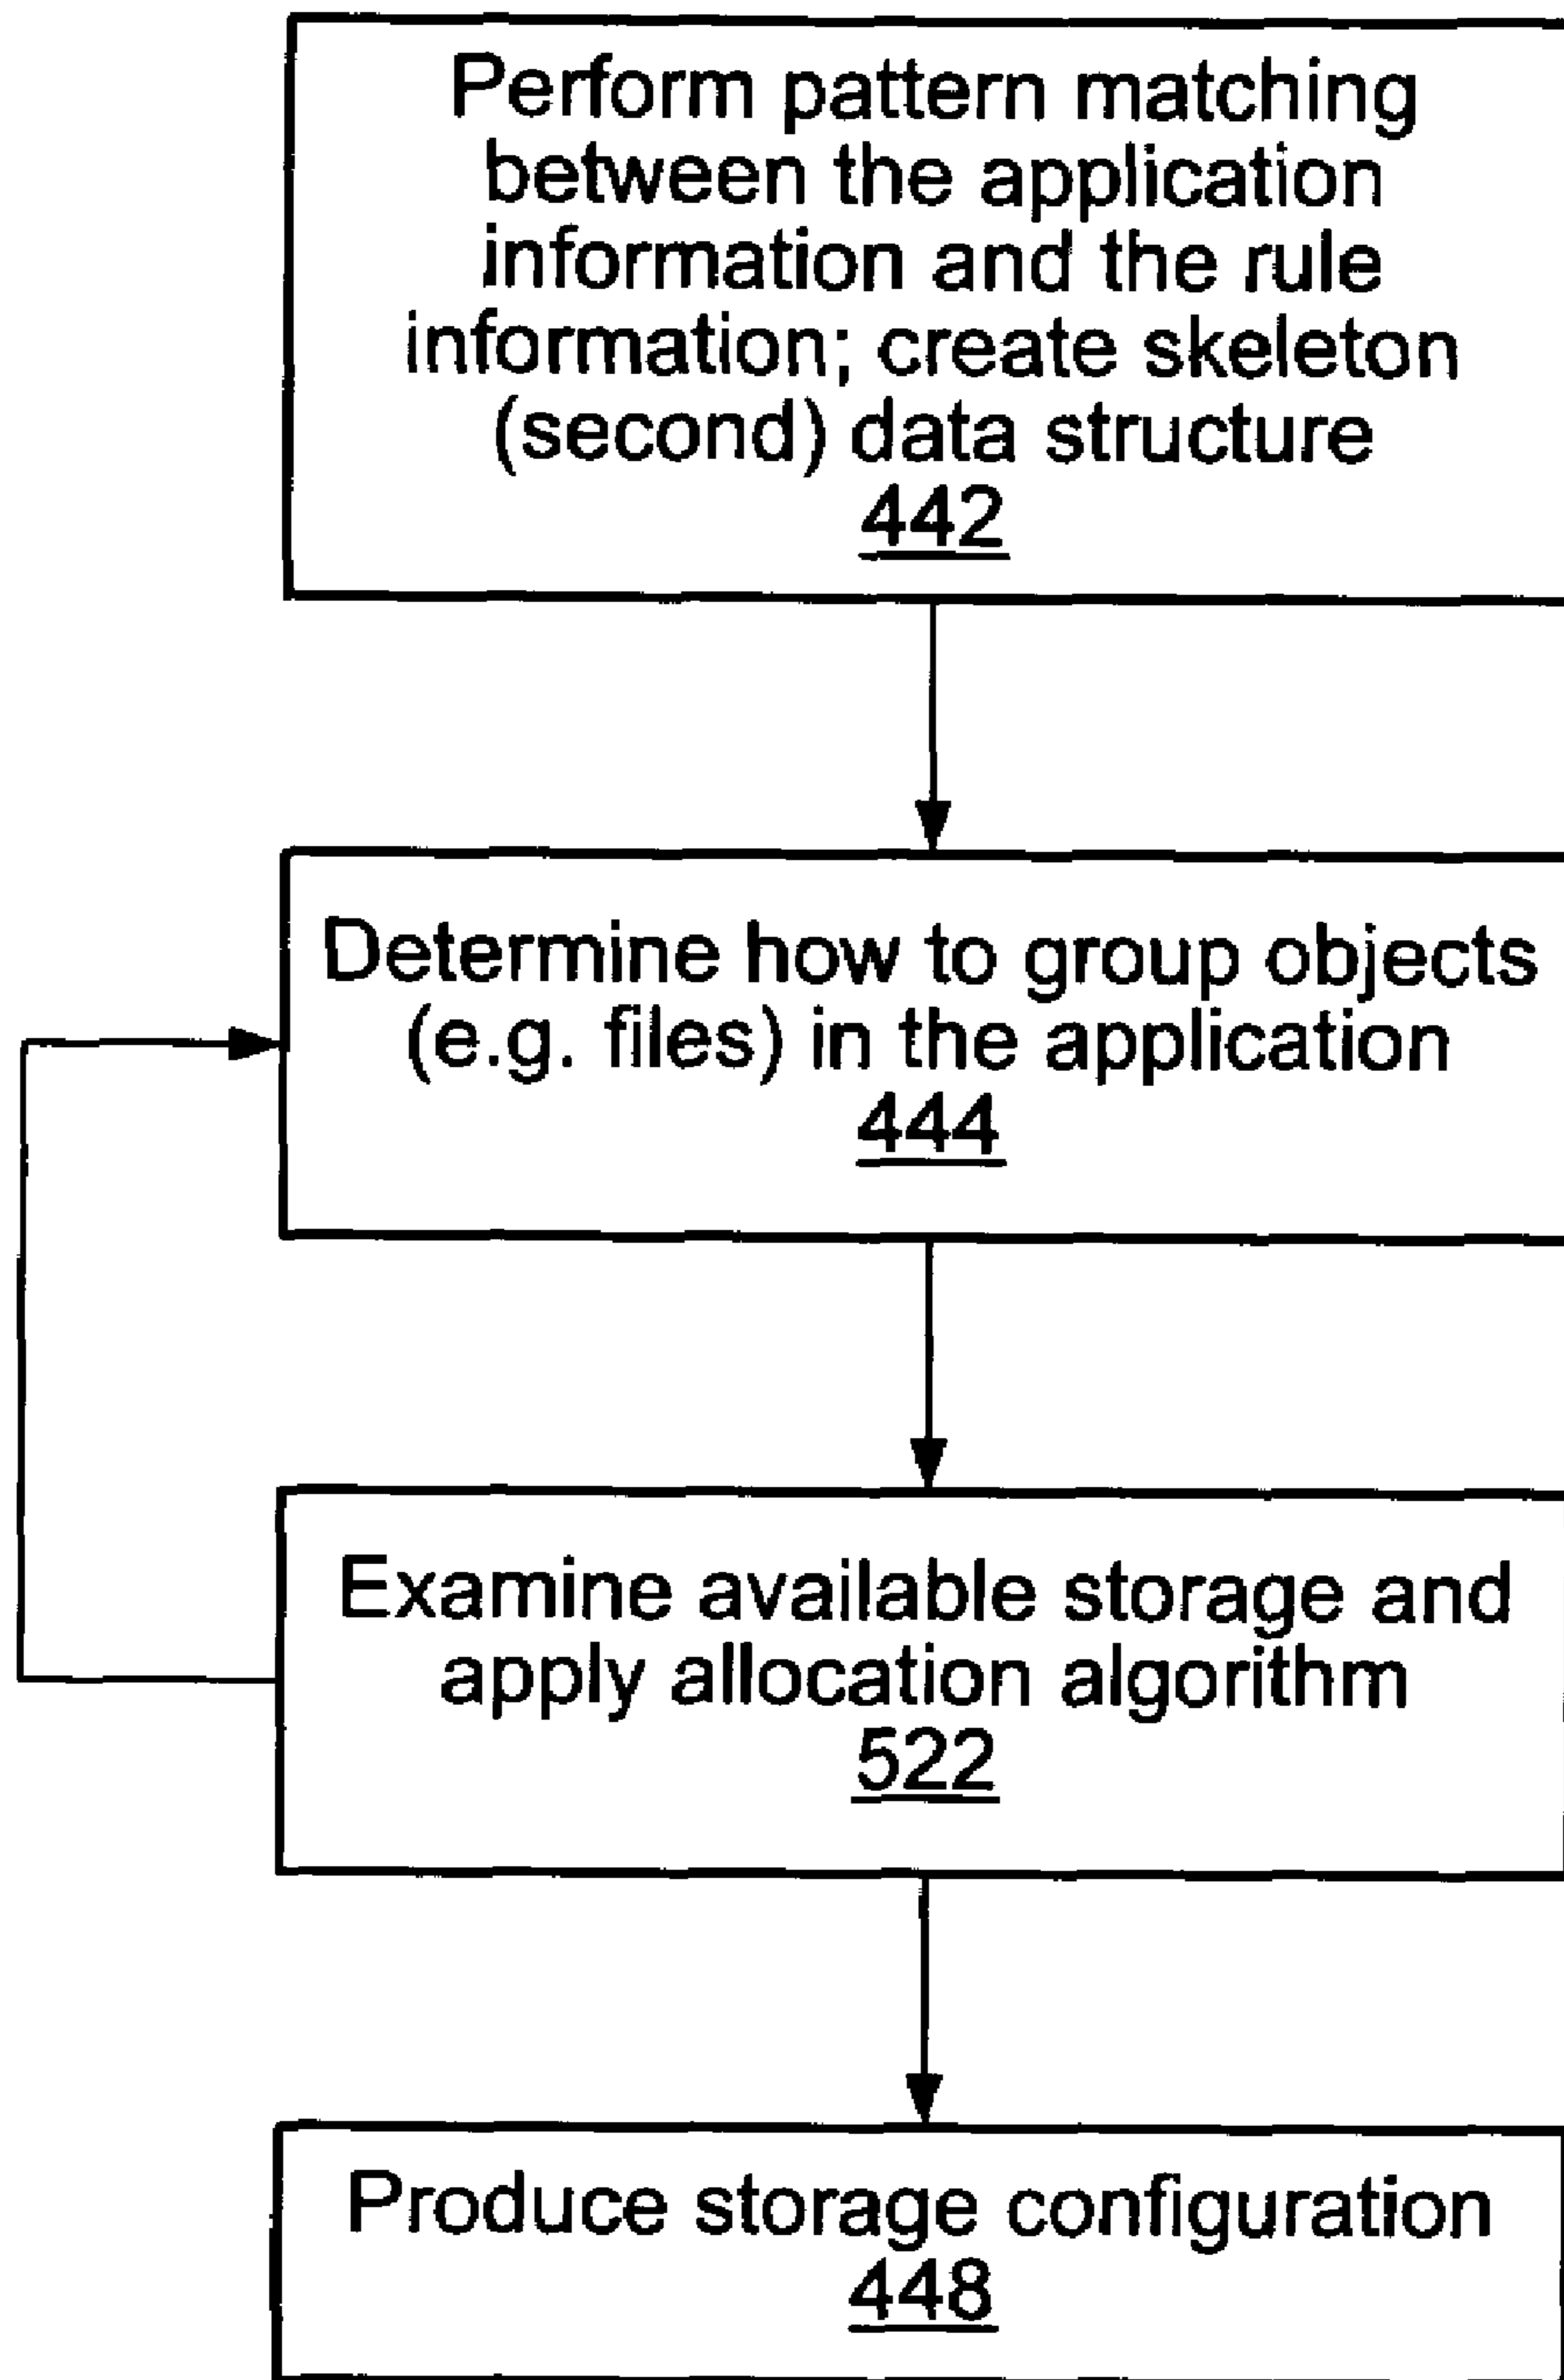
Fig. 24

**Discoverer Flowchart**



*Fig. 25*





*Fig. 26*

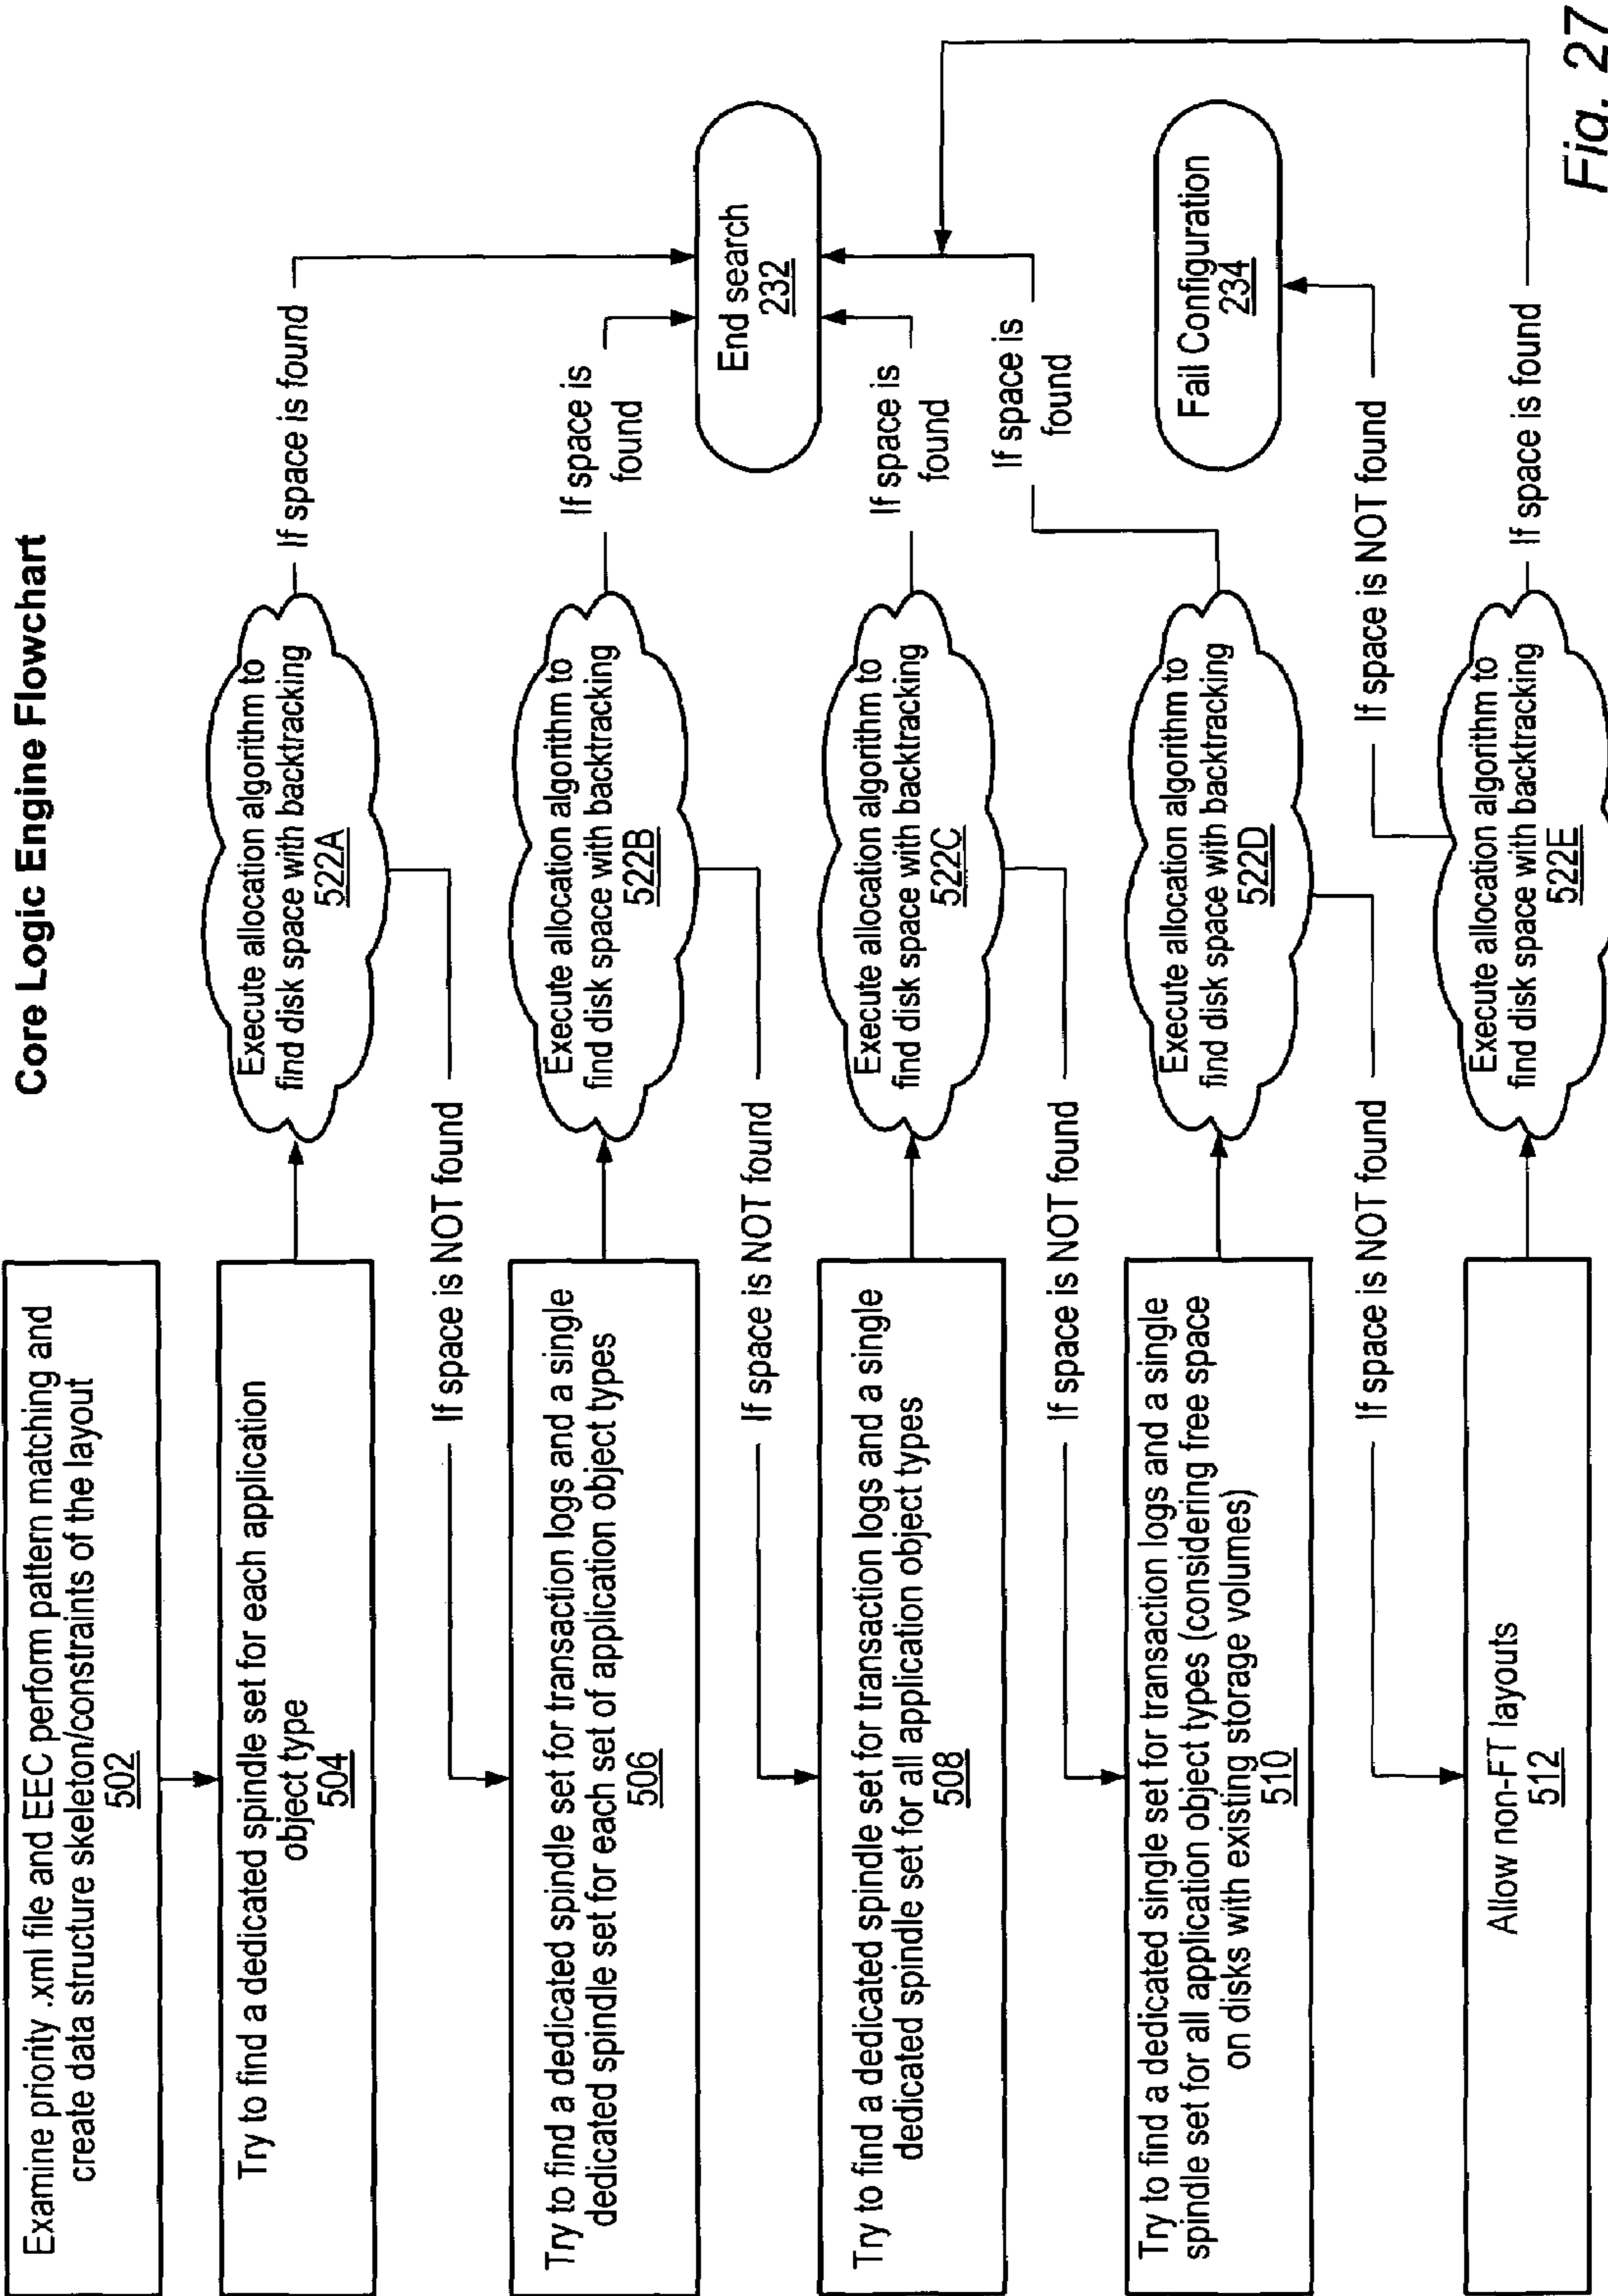


Fig. 27

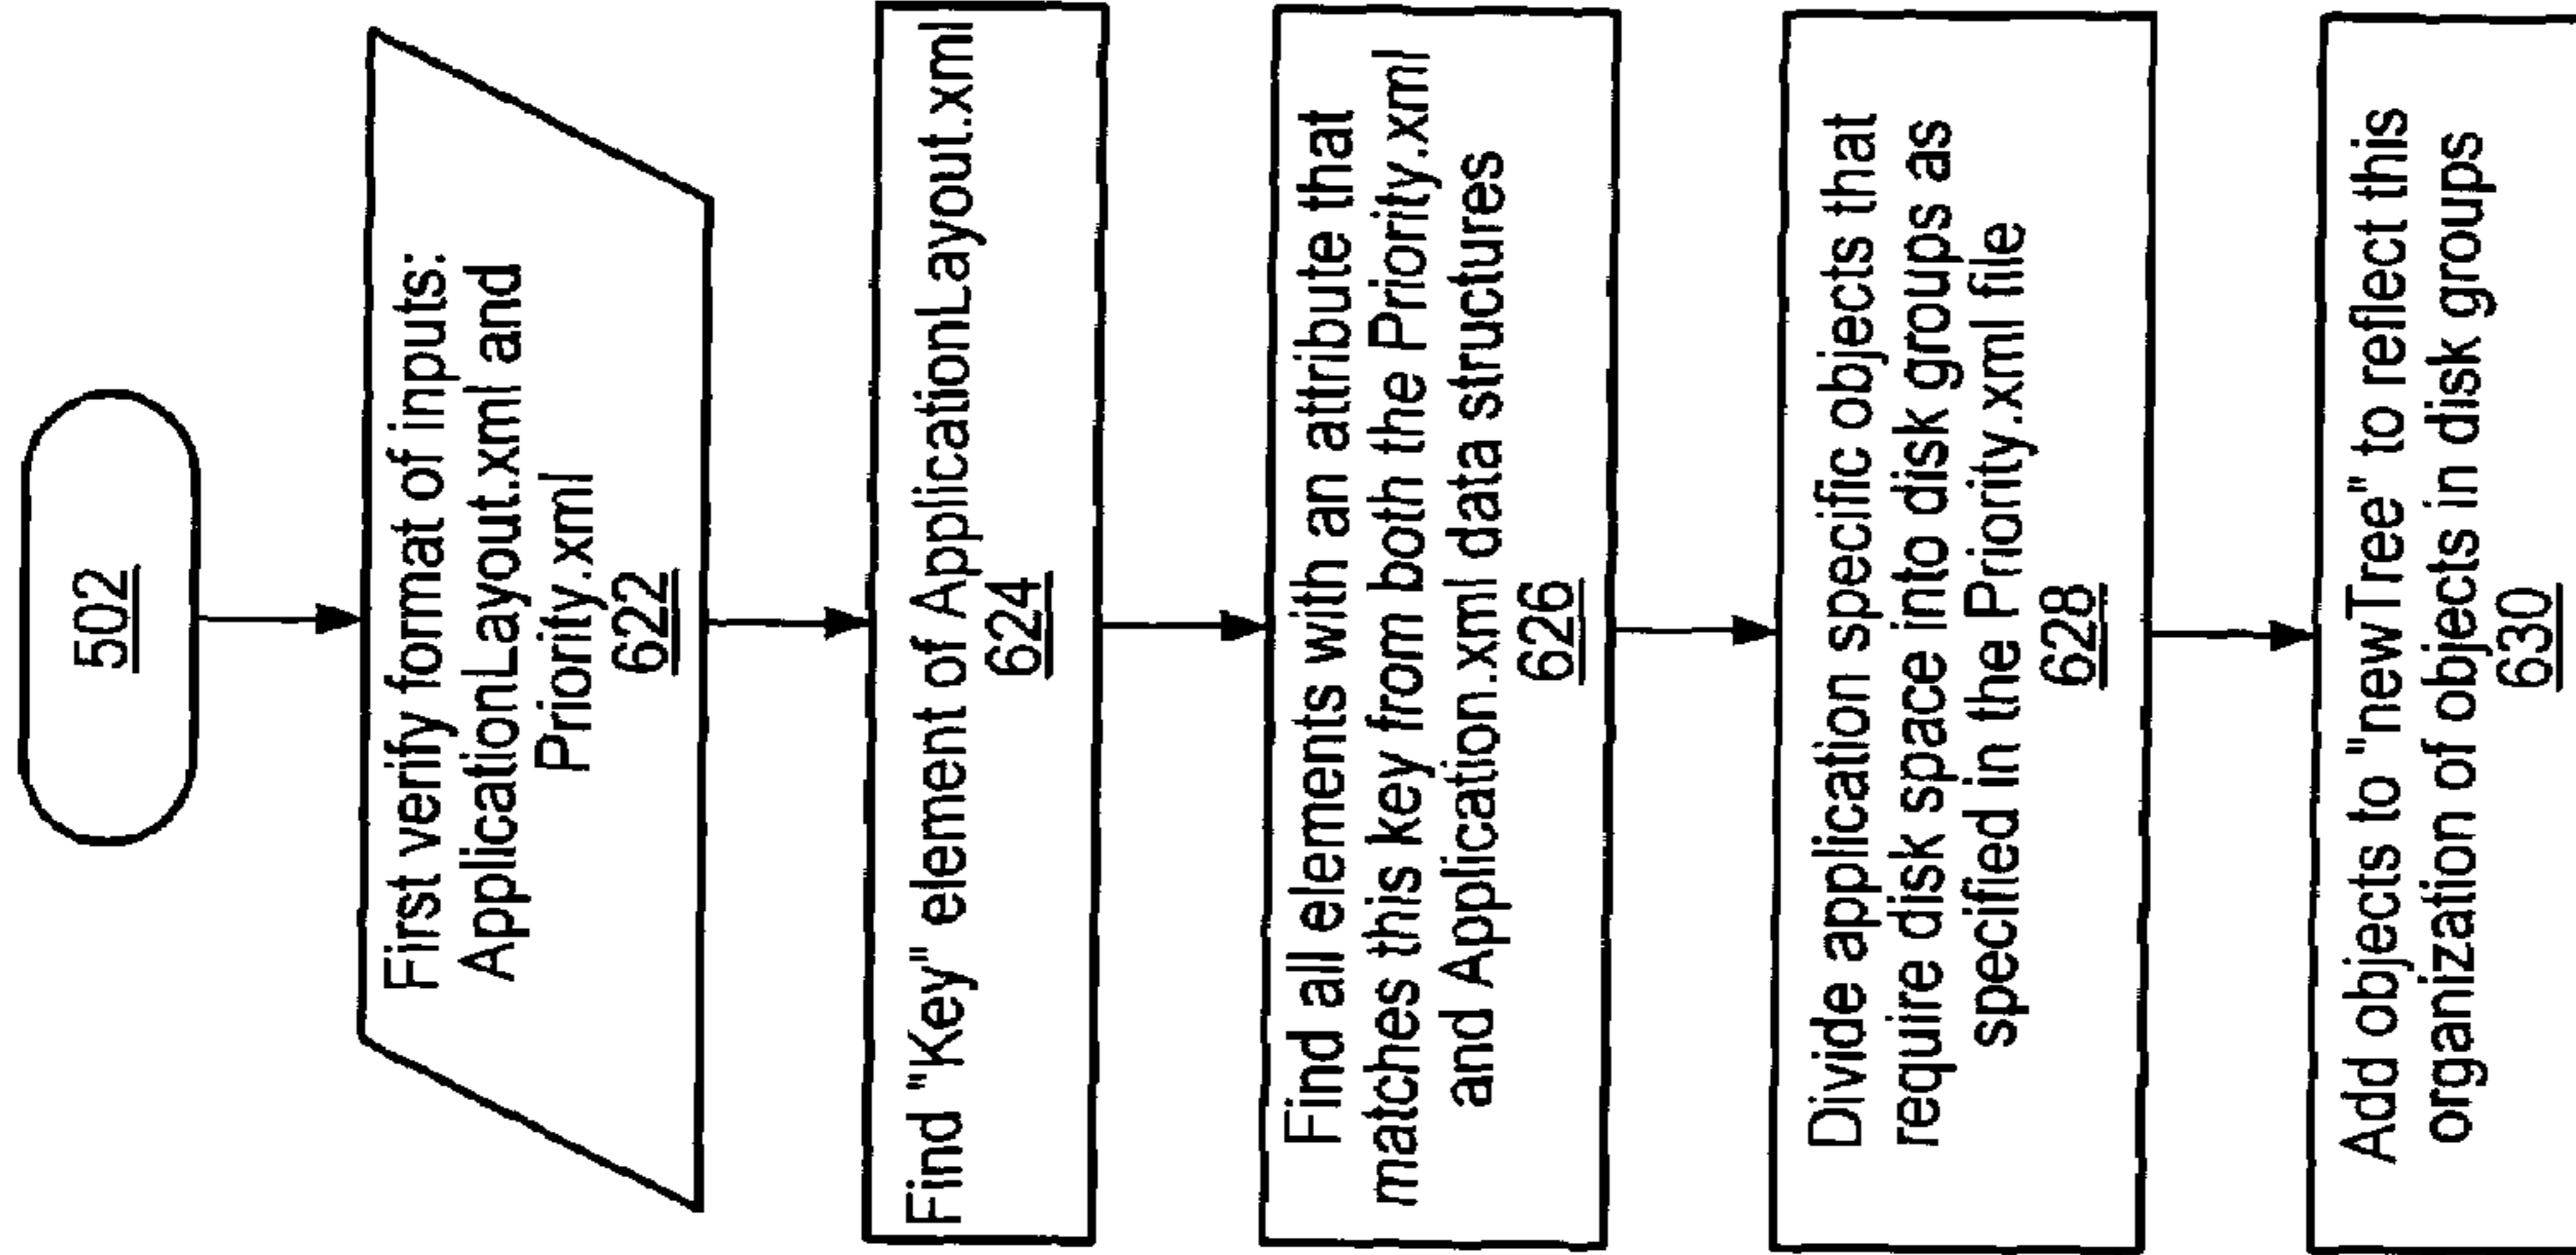


Fig. 29

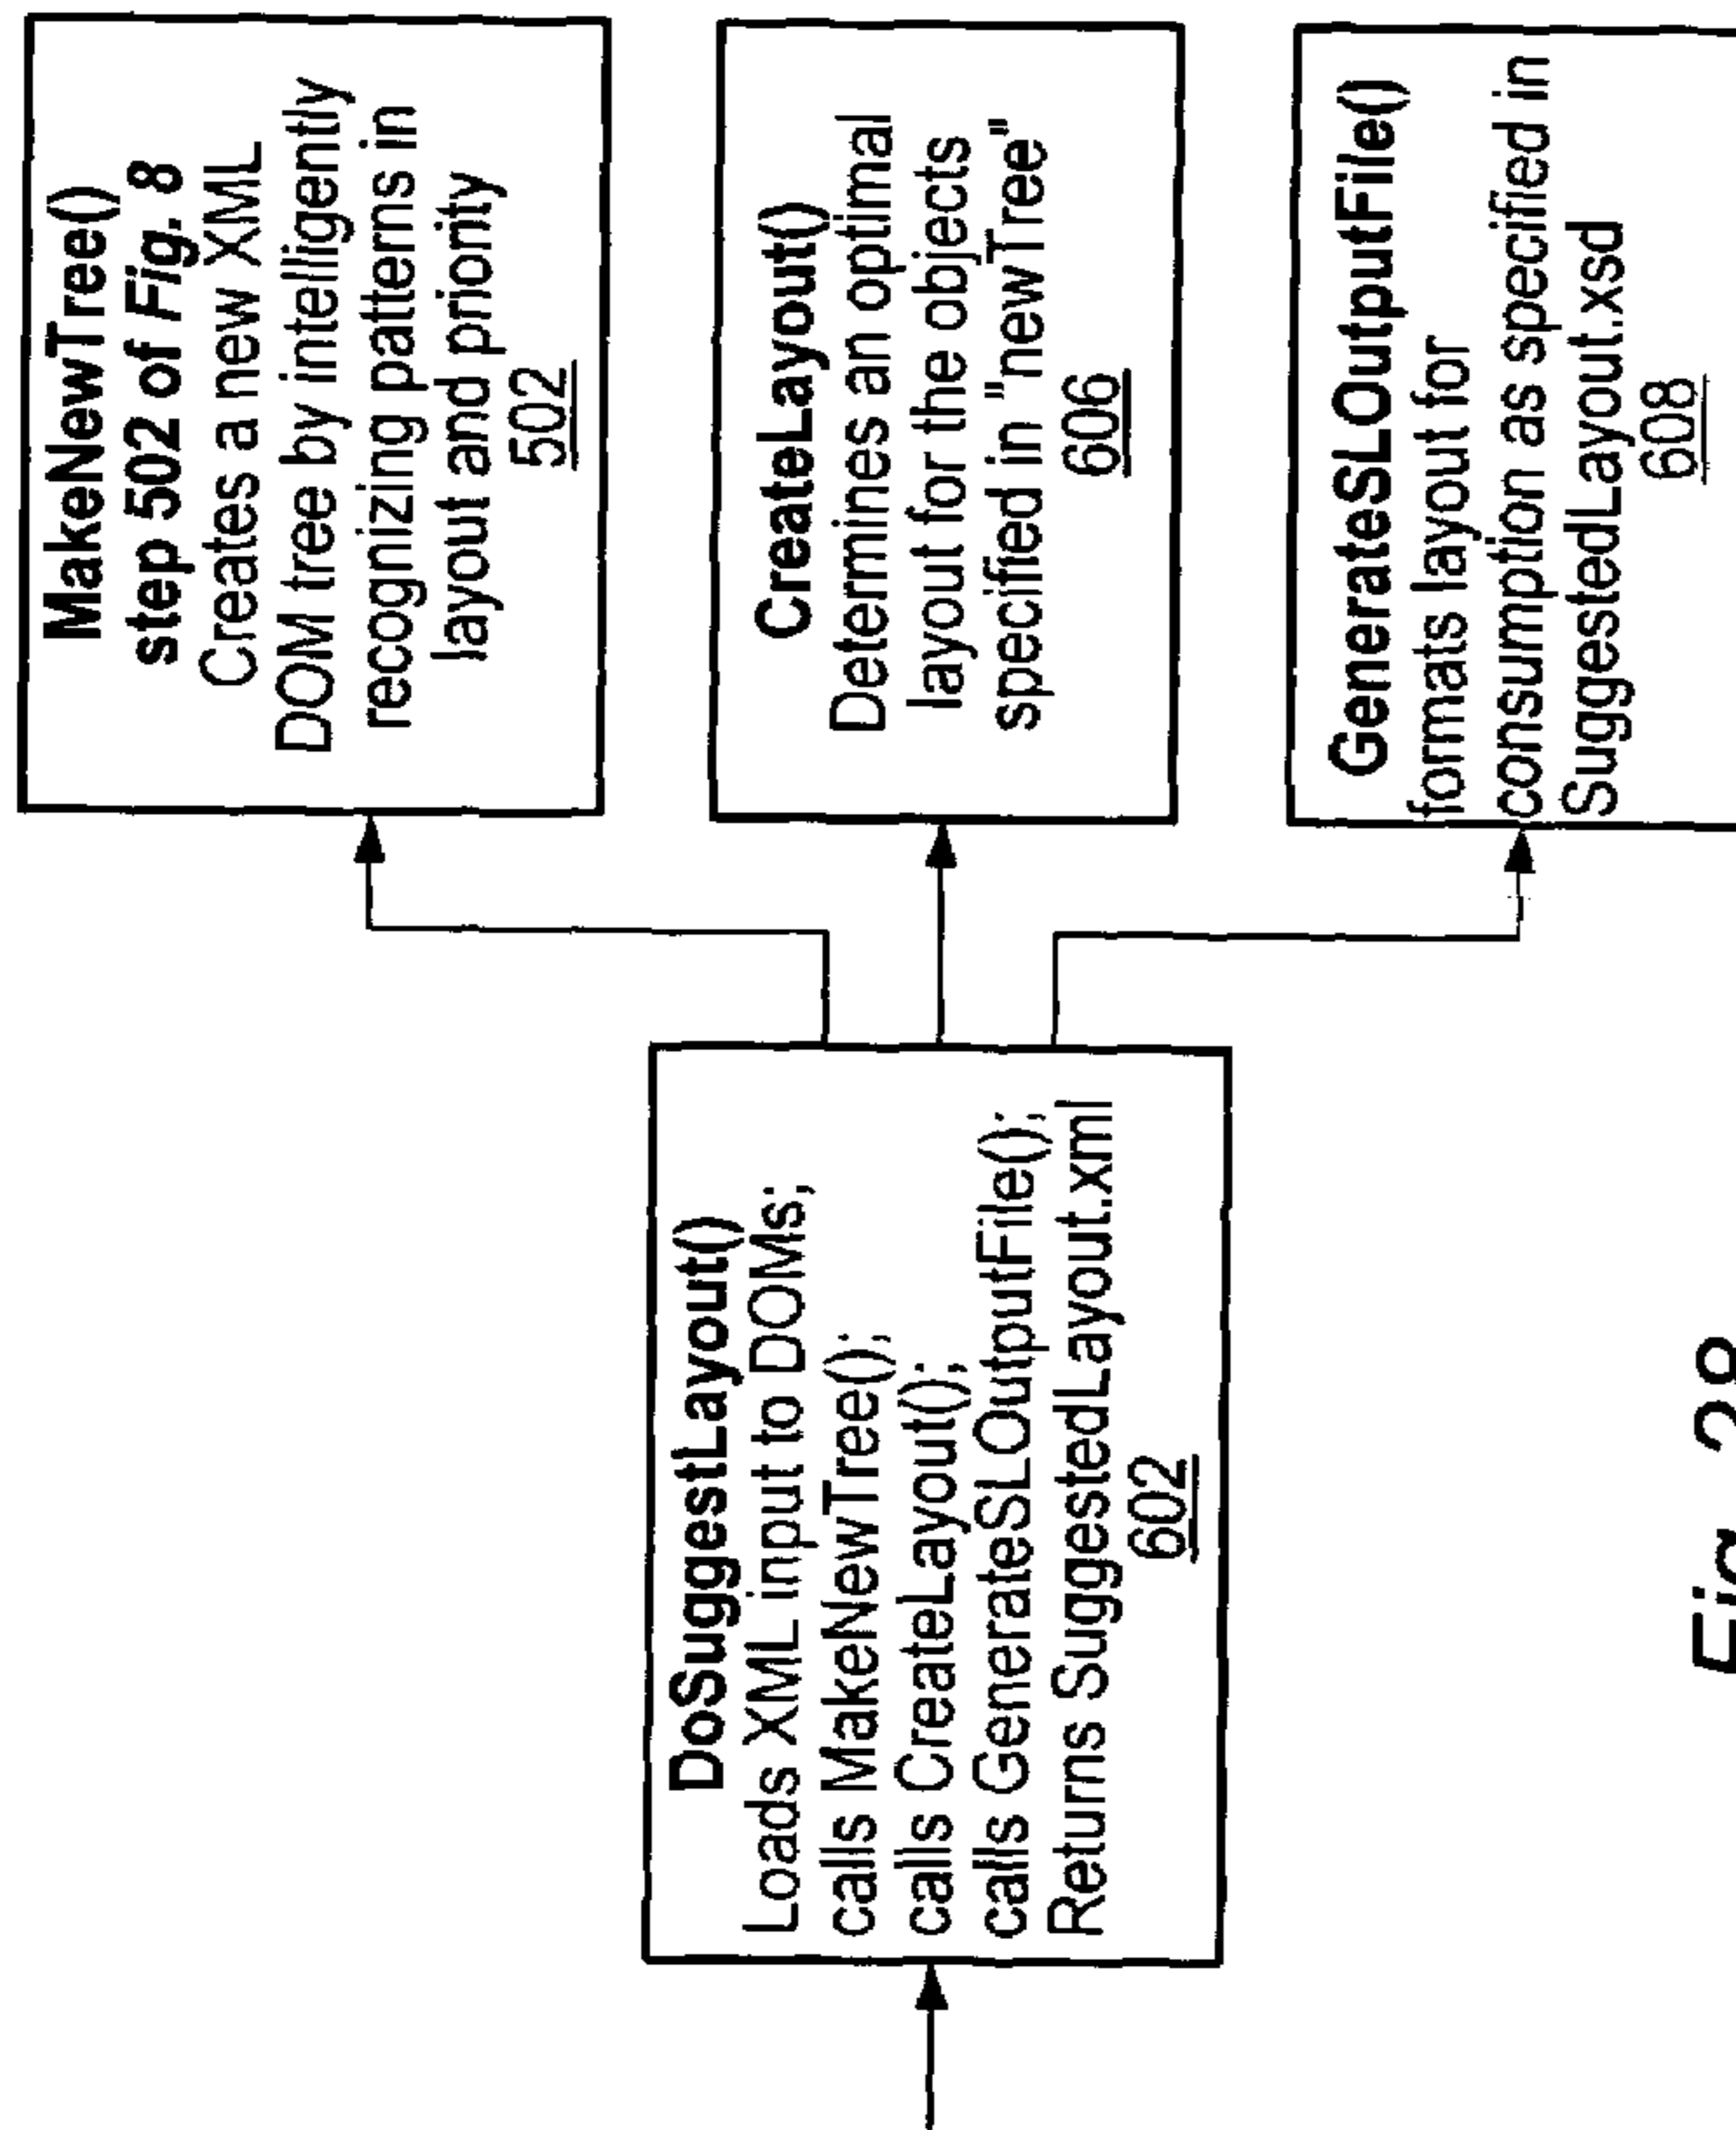


Fig. 28

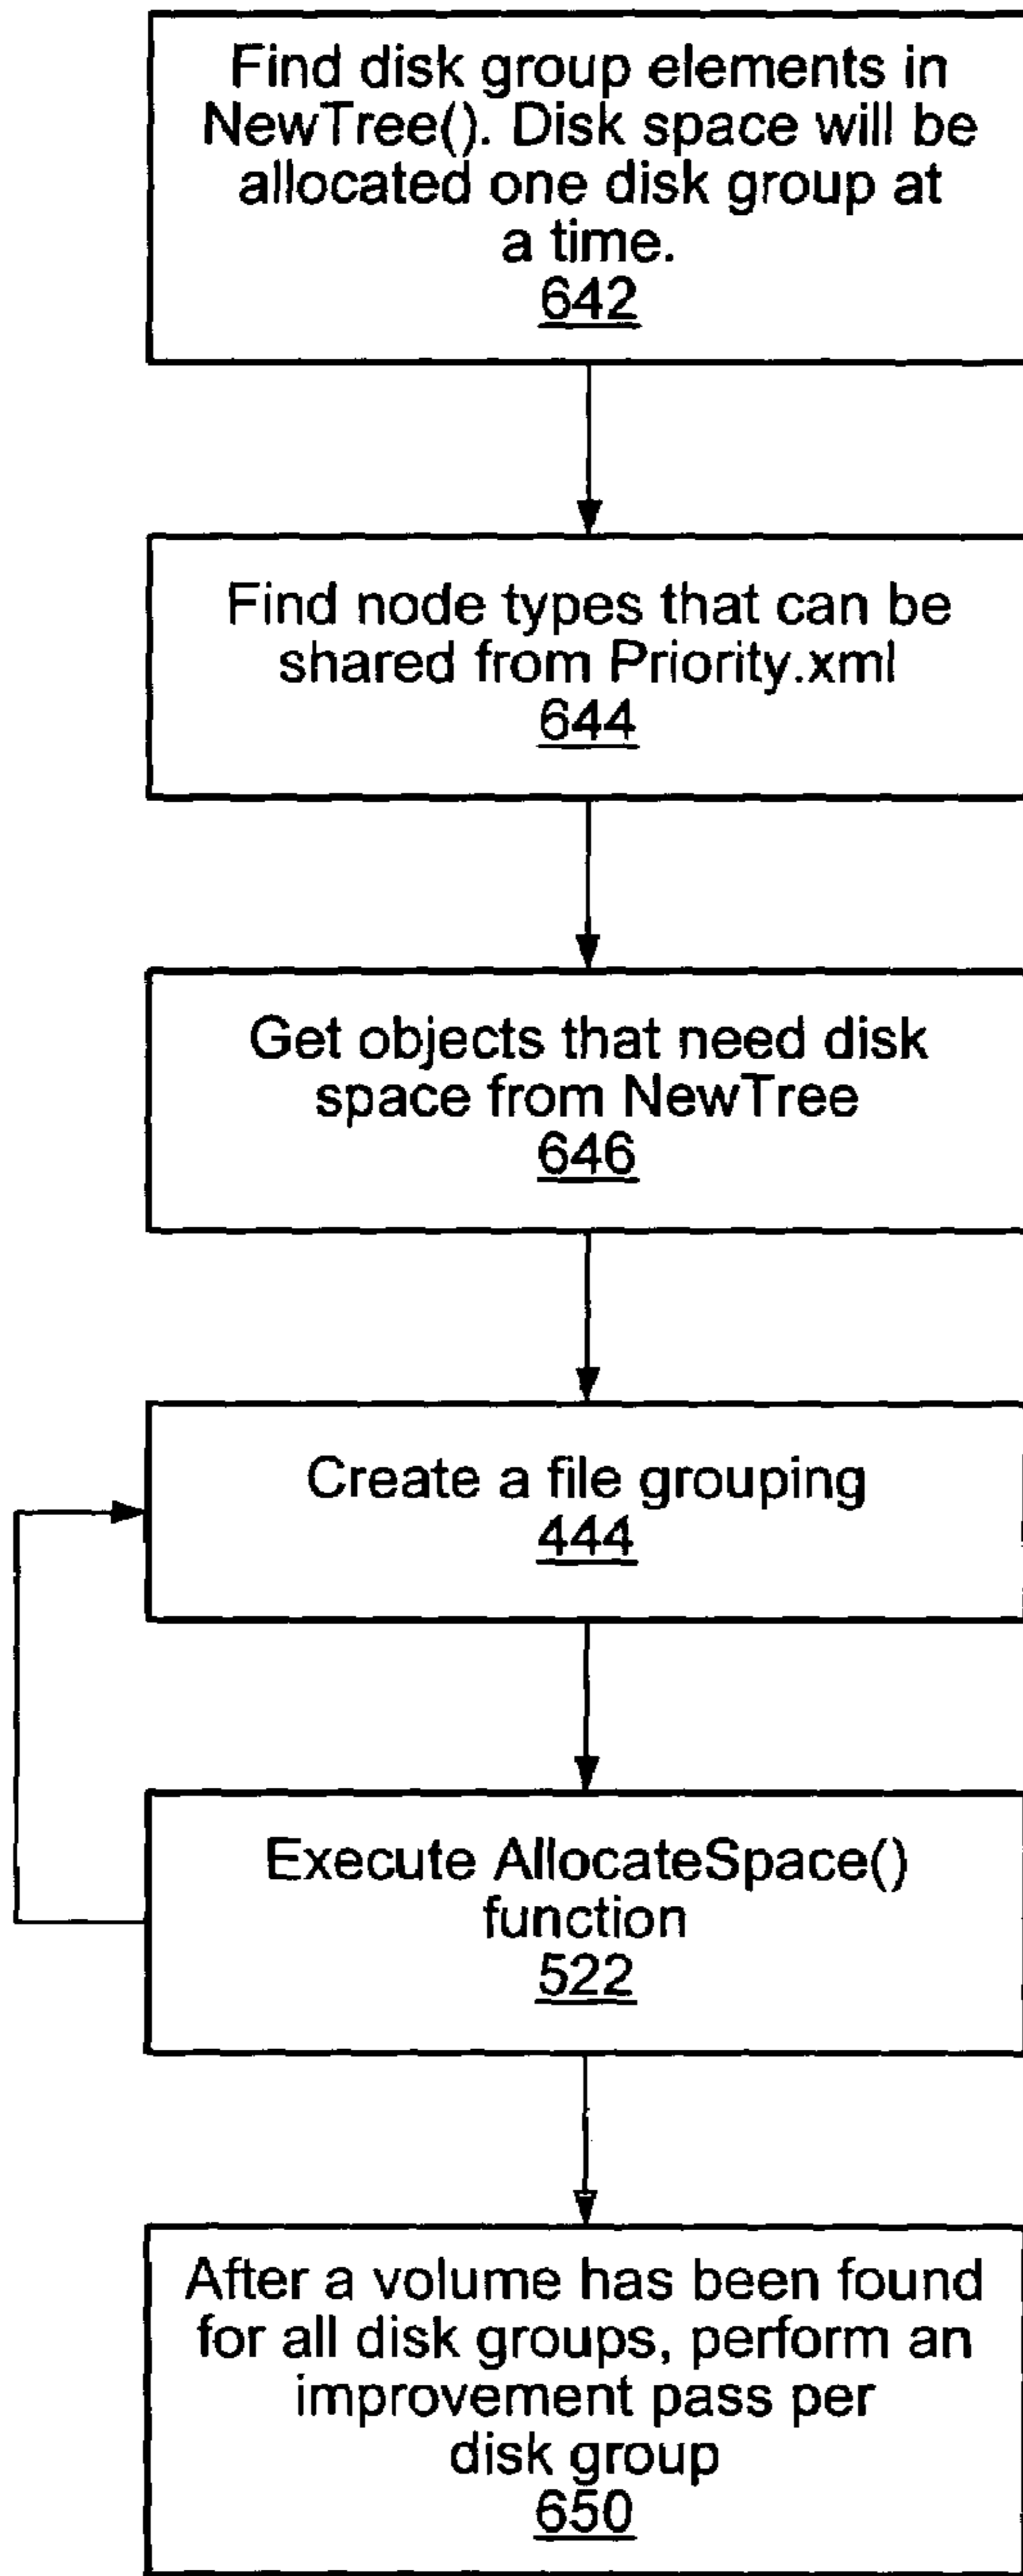


Fig. 30

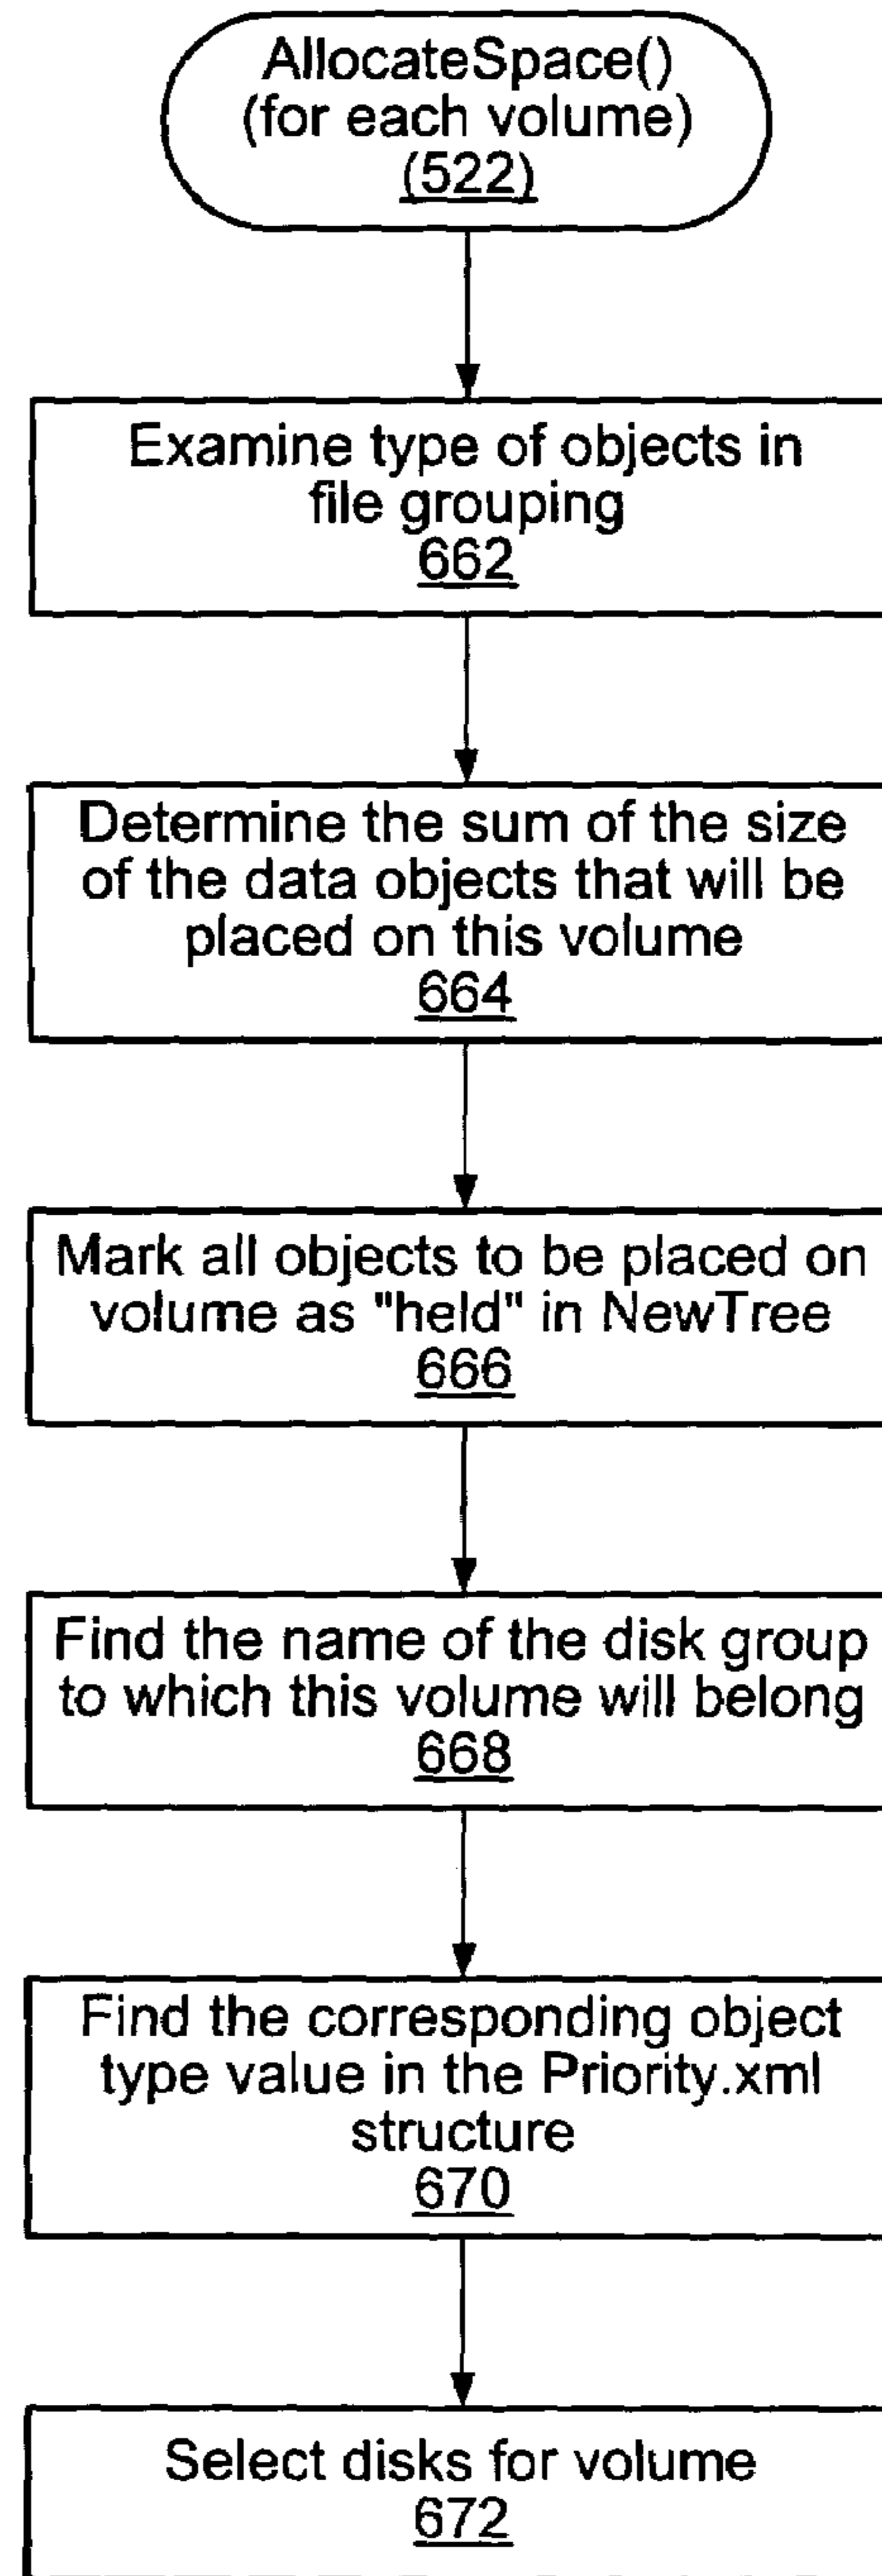


Fig. 31

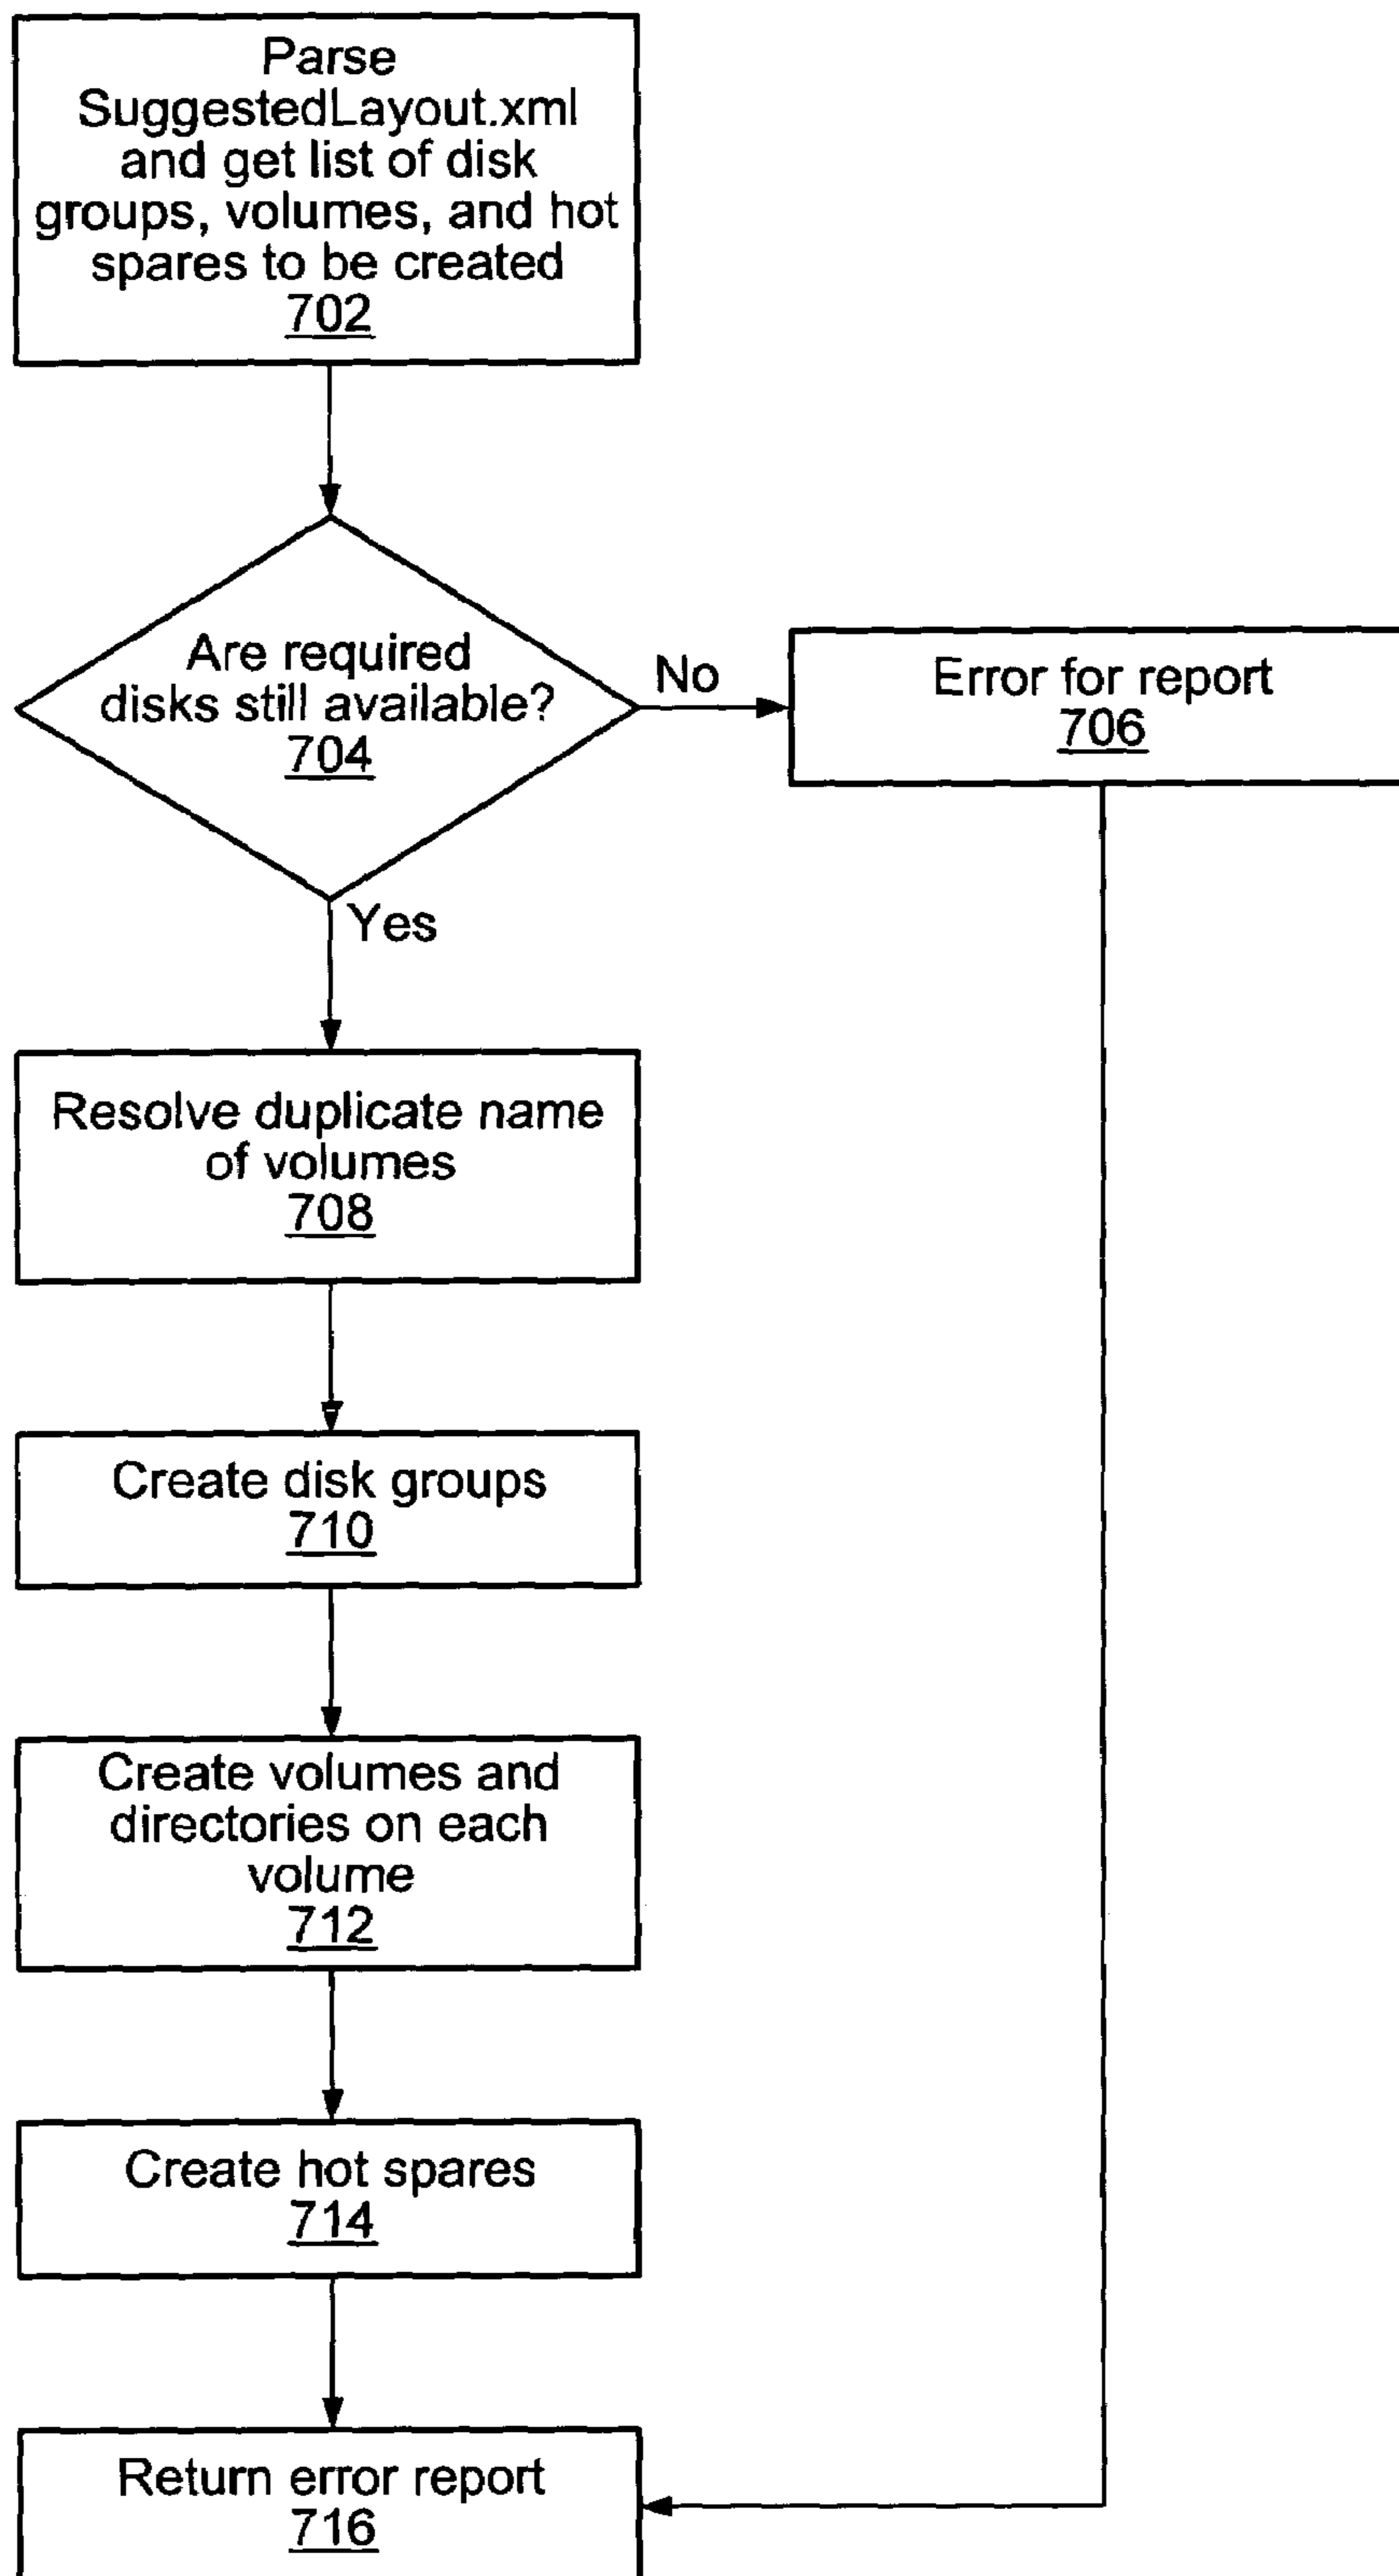


Fig. 32



**GRAPHICAL USER INTERFACE FOR  
CREATING A CUSTOMIZED STORAGE  
LAYOUT FOR AN APPLICATION**

PRIORITY CLAIM

This application is a continuation-in-part of U.S. patent application Ser. No. 10/327,561 titled "Storage Configurator for Determining an Optimal Storage Configuration for an Application" filed on Dec. 20, 2002 now U.S. Pat. No. 7,043,619, which claims benefit of priority of U.S. provisional application Ser. No. 60/348,870 titled "Storage Configurator" filed on Jan. 14, 2002.

BACKGROUND

1. Field of the Invention

This invention relates to a system and method for creating a customized storage configuration for an application, where the storage configuration specifies a storage layout for a plurality of data objects associated with the application.

2. Description of the Related Art

Network administrators are required to perform various types of tasks in setting up and running a network. One common task performed by a network administrator is creating a storage configuration, e.g., setting up and configuring storage on a server. Network administrators are required to create storage configurations for various applications, such as Microsoft Exchange, SQL databases, etc. An application may have a plurality of associated data objects, such as files or databases. Creating a storage configuration for the application may involve specifying the storage resources (e.g., disk groups, volumes, etc.) on which to store the various data objects.

The creation of a storage configuration for an application can be a difficult task. It would be desirable to provide a system operable to simplify this task by guiding the network administrator (user) through creation of the storage configuration. The system would preferably enable the user to specify desired properties of the storage configuration via a graphical user interface, validate the input specified by the user, and provide feedback regarding the specified storage configuration options.

SUMMARY OF THE INVENTION

Various embodiments of the invention relate to determining a storage configuration for an application, where the storage configuration specifies a storage layout for data objects associated with the application. A graphical user interface (GUI) for guiding the user through creation of the storage configuration may be displayed. The graphical user interface may be associated with a Storage Configurator software program. In one embodiment, the graphical user interface may comprise a series of panels or windows that guide the user through creation of the storage configuration in a wizard-like manner.

User input to specify properties of the storage configuration and/or properties of the application may be received to the graphical user interface. In various embodiments, any of various kinds of storage configuration or application properties may be specified, e.g., depending on the particular application for which the storage configuration is being created. As one example, the user input may include input specifying desired storage locations for one or more of the data objects associated with the application. For example, the user may specify a desired mapping of one or more data

objects to one or more storage resources (e.g., volumes or disk groups) of one or more storage devices (e.g., disks).

The Storage Configurator program may validate the user input and provide feedback to the user as appropriate. For example, if the user specifies a desired mapping of data objects to storage resources, the Storage Configurator program may determine whether the specified mapping of data objects to storage resources is valid. For example, determining whether the mapping is valid may comprise determining whether the storage resources specified by the mapping have sufficient storage capacity to hold the specified data objects. Determining whether the mapping is valid may also comprise validating the mapping against best practices of the application. If the Storage Configurator program determines that the mapping specified by the user is not valid or is inefficient, then information indicating the problem may be displayed on the graphical user interface.

The Storage Configurator program may determine a storage configuration for the application, based on the user input received to the graphical user interface. The storage configuration may specify a storage layout for data objects associated with the application. If the user specifies a desired mapping for one or more data objects to particular storage resources, then the specified mapping may be honored if possible. In one embodiment, the user may specify a desired storage layout for a subset of the application's data objects, and the Storage Configurator program may automatically determine a storage layout for the rest of the application's data objects. In another embodiment, the user may specify a desired storage layout for all of the application's data objects. In another embodiment, the Storage Configurator program may automatically determine a storage layout for all of the application's data objects.

In one embodiment, the Storage Configurator may provide an option allowing the user to save the storage configuration, e.g., in a file or database. Saving the storage configuration may allow the user to utilize the storage configuration at a later time. For example, the storage configuration may later be re-opened and modified.

In one embodiment, the Storage Configurator may also provide an option allowing the user to implement the storage configuration. Implementing the storage configuration may comprise storing the plurality of data objects associated with the application on one or more storage devices according to the storage layout specified by the storage configuration. For example, the user may simply click a button to cause the data objects to be stored according to the storage layout.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an exemplary network system according to one embodiment;

FIG. 2 is a flowchart diagram illustrating one embodiment of a method for determining and implementing a storage configuration for an application;

FIGS. 3-20 illustrate an exemplary graphical user interface (GUI) for determining and implementing a storage configuration for an application, according to one embodiment;

FIG. 21 is a block diagram illustrating a software architecture of one embodiment of the invention;

FIG. 22 is a block diagram illustrating the Storage Configurator Provider of FIG. 2 and related files;

FIG. 23 is a flowchart diagram illustrating one embodiment of a method of operation of a Storage Configurator software program to automatically determine an optimal storage configuration for an application;



FIG. 24 is a flowchart diagram illustrating an alternate embodiment of the method described in the flowchart of FIG. 4, where the user may provide constraints to the system;

FIG. 25 is a flowchart diagram illustrating automatic detection of existing storage in the system according to one embodiment;

FIG. 26 is a flowchart diagram illustrating automatic determination of a storage configuration according to one embodiment;

FIG. 27 is a more detailed flowchart diagram illustrating automatic determination of a storage configuration according to one embodiment;

FIGS. 28-31 are flowchart diagrams illustrating automatic determination of a storage configuration according to one embodiment; and

FIG. 32 is a flowchart diagram illustrating implementing a storage configuration on a storage system.

While the invention is described herein by way of example for several embodiments and illustrative drawings, those skilled in the art will recognize that the invention is not limited to the embodiments or drawings described. It should be understood that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the invention as defined by the appended claims.

#### DETAILED DESCRIPTION OF EMBODIMENTS

##### FIG. 1—Exemplary Network System

FIG. 1 illustrates an exemplary network system according to one embodiment. Embodiments of the invention may be used in any of various types of network or enterprise systems, and FIG. 1 illustrates one example of a network system. The network system may be a local area network (LAN), two or more interconnected local area networks, or a wide area network (WAN) comprising two or more distributed LANs coupled together, such as by the Internet, among other possible configurations.

As shown, the network system may include one or more client computer systems 102. The client computer systems 102 may store and/or execute various applications, one example being an electronic mail (email) or electronic contact/scheduling client program, such as Microsoft Outlook. The client computer systems 102 may execute other applications, such as programs for Internet access, database programs, and others. The one or more client computer systems 102 may be connected through a firewall 104 to a network device 106, such as a router, hub, or bridge. The client computer systems 102 may couple through the firewall 104 and/or network device 106 to one or more Internet servers 110. The Internet servers 110 may be connected to various other computer systems and/or storage devices, such as a WINS server 112, a backup server 114, and/or a storage area network (SAN) 116. These servers may in turn be connected to other servers which host other applications, such as Microsoft Exchange server 122 or an NT 4.0 PDC (Primary Domain Controller) 124, among others.

For many network systems, a network administrator or other user may operate to configure storage of one or more computer systems for various applications. For example, in the network system embodiment illustrated in FIG. 1, the user may configure storage of the server computer 122 which hosts the Microsoft Exchange application. In other

embodiments, the user may be required to configure or create a storage configuration for other applications, such as SQL databases, etc. Creating a storage configuration for an application may involve specifying the storage resources (e.g., disk groups, volumes, etc.) on which to store various data objects associated with the application, such as files or databases. For example, the Microsoft Exchange application utilizes various mailbox databases, among other data objects.

As noted above, creation of a storage configuration in the prior art can be a complicated and difficult task and typically requires significant expertise on the part of the network administrator. One embodiment of the present invention comprises a software program, referred to herein as the Storage Configurator program, executable on a computer system which operates to guide the network administrator or other user through creation of a storage configuration for a given application utilizing a storage system. As used herein, the term "storage system" includes any of various types of systems which provide storage capabilities, including, but not limited to, a server computer (e.g., having one or more RAID arrays), two or more coupled server computers, a storage area network (SAN), and other types of storage or memory systems. Such a software program may greatly simplify the task of the user in determining and creating storage configurations or storage layouts for a network system.

In various embodiments, the Storage Configurator software program may be stored in various locations, such as on a memory medium of one of the client computer systems 102, or a memory medium of one of the server computers. The Storage Configurator software program may also execute on any of various of the computer systems. In one embodiment, the Storage Configurator software program may be stored in a client computer system 102 used by the network administrator for system configuration and management.

The term "memory medium" is intended to include any of various types of memory devices for storing programs and/or data. For example, the term "memory medium" is intended to include an installation medium, e.g., a CD-ROM, floppy disks 104, or tape device; a computer system memory or random access memory such as DRAM, SRAM, EDO RAM, Rambus RAM, etc.; or a non-volatile memory such as a magnetic media, e.g., a hard drive, or optical storage. The memory medium may comprise other types of memory as well, or combinations thereof. In addition, the memory medium may be located in a first computer in which the programs are executed, or may be located in a second different computer which connects to the first computer over a network, such as the Internet. In the latter instance, the second computer may provide program instructions to the first computer for execution.

The software programs which implement embodiments of the present invention may be stored in a memory medium of one of the computers shown in FIG. 1, or in a memory medium of another computer, and executed by one or more CPUs. One or more CPUs executing code and data from a memory medium thus may comprise a means for performing the methods described herein. For example, a CPU executing code and data from a memory medium may comprise a means for guiding the user through creation of a storage configuration and/or implementing the storage configuration according to the methods described herein.

In general, the term computer system as used herein is defined to encompass any device having a processor which executes instructions from a memory medium. In different



## 5

embodiments, a “computer system” may take various forms, including a personal computer system, desktop computer, mainframe computer system, server computer system, another suitable device, or combinations thereof. A computer system may be attached to a network as part of a distributed computing environment.

FIG. 2—Method for Determining and Implementing a Storage Configuration

FIG. 2 is a flowchart diagram illustrating one embodiment of a method for determining and implementing a storage configuration for an application. It is noted that FIG. 2 illustrates a representative embodiment, and alternative embodiments are contemplated. Also, various elements may be combined, omitted, or performed in different orders.

In 10, a graphical user interface (GUI) for guiding the user through creation of the storage configuration may be displayed. The GUI may be associated with a Storage Configurator software program such as described above. For example, the GUI may be displayed in response to the user launching the Storage Configurator program. In various embodiments, the GUI may have any kind of appearance and may be presented in any of various ways. In one embodiment, the GUI may comprise a series of panels or windows that guide the user through creation of the storage configuration in a wizard-like manner.

In one embodiment, the GUI may present options allowing the user to step through the storage configuration creation process in either a Custom mode or a Typical mode. In the Typical mode, the user may be required to provide only a minimal amount of user input, and the storage configuration for the application may in large part be determined automatically. A method for automatically determining at least a portion of the storage configuration is described below. On the other hand, if the user has a good understanding of storage elements such as volumes and disk groups, the user can select the Custom option to manually specify various aspects of the storage configuration.

In 12, user input to specify properties of the storage configuration and/or properties of the application may be received to the graphical user interface. In various embodiments, any of various kinds of storage configuration or application properties may be specified, e.g., depending on the particular application for which the storage configuration is being created. The user input received in 12 may depend on whether the user chose the Custom mode or the Typical mode for the storage configuration creation process. In the Custom mode, the user may provide more specific input regarding the desired storage configuration.

As one example, the user input received in 12 may include input specifying names for one or more of the data objects (e.g., files or databases) associated with the application. The user may also specify desired storage locations for one or more of the data objects associated with the application. For example, the user may specify a desired mapping of one or more data objects to one or more storage resources (e.g., volumes or disk groups) of one or more storage devices (e.g., disks).

The Storage Configurator program may validate the user input received in 12 and provide feedback to the user as appropriate. For example, if the user specifies a desired mapping of data objects to storage resources, the Storage Configurator program may determine whether the specified mapping of data objects to storage resources is valid. For example, determining whether the mapping is valid may comprise determining whether the storage resources speci-

## 6

fied by the mapping have sufficient storage capacity to hold the specified data objects. Determining whether the mapping is valid may also comprise validating the mapping against best practices of the application. For example, the application may have an associated set of rules specifying optimal locations for particular data objects. If the Storage Configurator program determines that the mapping specified by the user is not valid or is inefficient, then information indicating the problem may be displayed on the graphical user interface.

In 14, the Storage Configurator program may determine a storage configuration for the application, based on the user input received in 12. The storage configuration may specify a storage layout for data objects associated with the application. If the user specified a desired mapping for one or more data objects to particular storage resources (e.g., in the Custom mode), then the specified mapping may be honored if possible. In one embodiment, the user may specify a desired storage layout for a subset of the application’s data objects, and the Storage Configurator program may automatically determine a storage layout for the rest of the application’s data objects. In another embodiment, the user may specify a desired storage layout for all of the application’s data objects. In another embodiment, the Storage Configurator program may automatically determine a storage layout for all of the application’s data objects (e.g., in the Typical mode).

In 15, the storage configuration may be displayed on the GUI. In various embodiments, the storage configuration may be displayed in various ways. In one embodiment, the Storage Configurator program may be operable to display the storage configuration using different views selectable by the user. For example, the user may select a “Storage Group” view to cause the Storage Configurator program to display the storage configuration based on various storage groups. As another example, the user may select a “Volume” view to cause the Storage Configurator program to display the storage configuration based on volumes utilized in the storage configuration.

In one embodiment, the GUI may provide an option allowing the user to save the storage configuration, e.g., in a file or database. If the user wants to save the storage configuration, then the storage configuration may be saved or stored in 16. In one embodiment the storage configuration may be saved as an XML file. Saving the storage configuration may allow the user to utilize the storage configuration at a later time. For example, the storage configuration may later be re-opened, as indicated in 18. This may comprise retrieving the stored storage configuration and re-displaying the storage configuration in the GUI. As indicated in 20, the user may optionally modify properties of the previously saved storage configuration.

In one embodiment, the GUI may also provide an option allowing the user to implement the storage configuration. If the user wants to implement the storage configuration, then the storage configuration may be implemented in 22. Implementing the storage configuration may comprise storing the plurality of data objects associated with the application on one or more storage devices according to the storage layout specified by the storage configuration. For example, the user may simply click a button to cause the data objects to be stored according to the storage layout.

FIGS. 3-20: Exemplary Graphical User Interface

FIGS. 3-20 illustrate an exemplary graphical user interface (GUI) for determining and implementing a storage



configuration for an application, according to one embodiment. This exemplary GUI is useable to create a storage configuration for the Microsoft Exchange Server application in conjunction with Veritas Volume Manager. However, alternative embodiments may comprise GUI's for creating storage configurations for any of various other applications in conjunction with any of various other storage management systems. The exemplary GUI of FIGS. 3-20 includes a series of screens or panels that are displayed in sequence to guide the user in a wizard-like manner through the process of creating the storage configuration for Microsoft Exchange Server.

The GUI may be launched in response to various conditions or actions. For example, the user may select a menu item, click a toolbar button, or perform other kinds of actions to request the GUI to be launched. FIG. 3 illustrates a welcome panel that the user sees when the GUI is launched. The welcome panel provides a brief description of the goals of the GUI.

After clicking Next, the panel shown in FIG. 4 appears. This panel provides two options for the user based on the user's experience. The user can select the Typical option if the user wants a storage configuration to automatically be determined. However, if the user has a good understanding of storage elements such as volumes and disk groups, the user can select the Custom option to manually specify the storage configuration.

FIG. 5 illustrates a panel that is displayed if the user chooses the Typical option. This panel provides a description of the tasks the user will be guided through for creating the storage configuration in the Typical mode. As shown, these tasks include identifying all the useable storage groups useable for database creation, specifying information about the storage groups, specifying information about databases within the storage groups, and selecting empty disks to be used.

FIG. 6 illustrates a panel that is displayed if the user chooses the Custom option. This panel provides a description of the tasks the user will be guided through for creating the storage configuration in the Custom mode. As shown, these tasks include identifying all the useable storage groups useable for database creation, specifying information about the storage groups, and specifying information about databases within the storage groups.

FIG. 7 illustrates a panel that is common to both the Typical and Custom modes. This panel allows the user to specify attributes of the storage groups that the user wants to configure or create. The various components shown on the panel are described as follows:

"Existing storage groups that were found on Basic Disk Group"—Indicates the number of existing Storage Groups that reside on Basic Disk Group.

"Move existing transactions logs and database that reside on Basic Disk Groups to Dynamic Disk Groups"—Allows the user to specify whether or not to move the existing Storage Groups that reside on Basic Disk Group to Dynamic Disk Groups. By selecting the checkbox, the Storage Group panel becomes editable, otherwise a read-only panel will be displayed.

"Existing Storage Groups that were found on Dynamic Disk Groups"—Indicates the number of existing Storage Groups that reside on Dynamic Disk Groups.

"Existing Storage Groups that were unconfigurable"—Indicates the number of existing Storage Groups that contains Exchange objects residing across different disk groups that the wizard cannot configure or that have already reached the maximum number of Exchange objects.

"Additional Storage Groups to be created"—Indicates the number of new Storage Groups that the user wants to create.

"Exchange Server type"—Indicates the type of the Exchange Server that the user is configuring. The Enterprise type can have a total of four Storage Groups, whereas the Standard type can have only one Storage Group on the Exchange server.

"Total number of Storage Groups"—Indicates the number of Storage Groups. For each Storage Group, a panel is presented for the user to specify attributes of the Storage Group.

"(# is configurable)"—Indicates the number of Storage Groups that are configurable. Configurable means the Storage Group is either new, or can still be added one or more mailbox databases, or moved from Basic Disk Group to Dynamic Disk Group.

FIG. 8 illustrates a panel for configuring storage groups that is displayed in Typical mode. This panel may be displayed N times, where N is the number of storage groups. The panel allows the user to specify attributes of each respective storage group. The various components shown on the panel are described as follows:

"Storage Group X of N"—Indicates the index of the Storage Group that the user is currently configuring (X) and the total number of Storage Groups (N).

"Storage Group name"—The name of the Storage Group. The field is grayed out if the Storage Group already exists or is unconfigurable. If the field is editable, a unique name is generated and used to populate the field when the panel is initially displayed. The user may modify the name as desired.

"Disk Group name"—The name of the Disk Group. This field is grayed out if the Storage Group is unconfigurable, or already exists and resides on a Dynamic Disk Group. For new Storage Groups and existing Storage Groups that reside on Basic Disk Group, a unique name is generated and used to populate the field. The user may modify the name as desired.

"Databases" table—A table view of all the mailbox databases within the Storage Group. For mailbox databases that already exist, the icons in the Name column are grayed out. These rows are also read-only with User Profile column filled in with "N/A". The Size column indicates the combined size of edb and stm database files. User Profile needs to be selected by the user during mailbox database creation. The table may be updated as the user utilizes the "New", "Edit", and "Delete" buttons, as described below.

"New" button—This button is enabled if the Storage Group is configurable and has not reached its maximum number of mailbox databases. In response to clicking this button, a panel appears (see FIG. 9) to allow the user to specify attributes of the mailbox database to be added. A corresponding row then appears in the "Databases" table.

"Edit" button—This button is enabled if one of the new mailbox databases is selected. In response to clicking this button, a panel appears (see FIG. 9) to allow the user to edit attributes of the selected mailbox database. The corresponding row in the "Databases" table may then be updated.

"Delete" button—This button is enabled if one of the new mailbox databases is selected. In response to clicking this button, the selected mailbox database is removed from the "Databases" table.

FIG. 9 illustrates a panel that is displayed to create a mailbox database (in Typical or Custom Mode) or edit a mailbox database (in Typical Mode). As shown, the panel allows the user to specify sizes of the EDB and STM database files and select a user profile for each one. The user



can also specify a name for the mailbox database and provide a name for the database file. The database file name is the file name for both mailbox and streaming files.

FIG. 10 illustrates a panel for configuring storage groups that is displayed in Custom mode. This panel may be displayed N times, where N is the number of storage groups. The panel allows the user to specify attributes of each respective storage group. The various components shown on the panel are described as follows:

“Storage Group X of N”—Indicates the index of the Storage Group that the user is currently configuring (X) and the total number of Storage Groups (N).

“Storage Group name”—The name of the Storage Group. The field is grayed out if the Storage Group already exists or is unconfigurable. If the field is editable, a unique name is generated and used to populate the field when the panel is initially displayed. The user may modify the name as desired.

“Disk Group name”—The name of the Disk Group. This field is grayed out if the Storage Group is unconfigurable, or already exists and resides on a Dynamic Disk Group. For new Storage Groups and existing Storage Groups that reside on Basic Disk Group, a unique name is generated and used to populate the field. The user may modify the name as desired.

“Databases” table—A table view of all the mailbox databases within the Storage Group. Each mailbox database is represented by two rows: edb and stm rows. A TransLog row gets created by default for new Storage Groups. Two rows get inserted with no volume assignment during new database creation. “Assigned Volume” indicates the volume that the mailbox database file belongs.

“Set Volume” button—This button is enabled if the Storage Group is configurable and a configurable row is currently selected. A configurable row is a row that contains a mailbox database file that is either newly created or its volume can be reassigned. In response to clicking this button, a panel appears (see FIG. 11) to allow the user to specify a volume.

“Hot Spare” button—This button is enabled if the Storage Group is configurable. In response to clicking this button, a panel appears (see FIG. 13) to allow the user to select one or more disks to set aside as hot spare.

“New” button—This button is enabled if the Storage Group is configurable and has not reached its maximum number of mailbox databases. In response to clicking this button, a panel appears (see FIG. 9) to allow the user to specify attributes of the mailbox database to be added. Two corresponding rows (edb row and stm row) then appear in the “Databases” table.

“Edit” button—This button is enabled if one of the new mailbox database files or TransLog object is selected. In response to clicking this button, a panel appears (see FIG. 14) to allow the user to edit attributes of the selected mailbox database file or TransLog object. The corresponding row in the “Databases” table may then be updated.

“Delete” button—This button is enabled if one of the newly created mailbox database files is selected. In response to clicking this button, the selected row and the corresponding edb or stm row is removed from the “Databases” table.

FIG. 11 illustrates a panel that is displayed in Custom mode to allow a user to select an existing volume or create and select a new volume. The various components shown on the panel are described as follows:

“Disk Group name”—Displays the disk group name.

“Available Volumes” table—Displays volumes that belong to the disk group, including their names, layouts, and

sizes. These may include volumes that already exist or newly created volumes. Existing volumes are displayed with a grayed out volume icon in the Name column.

“New” button—This button allows the user to create new volumes. In response to clicking this button, a panel appears (see FIG. 12) to allow the user to specify attributes of the volume to be added. A corresponding row then appears in the “Available Volumes” table. It is noted that the volumes are not actually created until the storage configuration is implemented.

“Edit” button—This button is enabled when a newly created volume is selected. In response to clicking this button, a panel appears (see FIG. 12) to allow the user to edit attributes of the selected volume.

“Delete” button—This button is enabled when a newly created volume is selected. In response to clicking this button, the selected volume is removed from the “Available Volumes” table.

“OK” button—This button becomes enabled when the user selects a volume in the “Available Volumes” table. Before the volume gets assigned to the selected mailbox database file, the volume size is validated to ensure that it has enough space to hold the file the user specified. A warning message is displayed if the volume size is not large enough to hold the file. The volume will not be assigned to the database file until a volume with the proper size is selected.

FIG. 12 illustrates a panel that is displayed in Custom mode to allow a user to create or edit a volume. The various components shown on the panel are described as follows:

“Name”—Name of the volume.

“Label”—Label of the volume.

“Size”—Size of the volume.

“Layout” fields—Used to set information regarding the volume layout. (Refer to Veritas Volume Manager documentation for information regarding volume layout.)

“Mirror” fields—Used to set information regarding mirroring for the volume. (Refer to Veritas Volume Manager documentation for information regarding mirroring.)

“Assign disks for Volume”—The user can move disks from the “Available disks” list to the “Selected disks” list by doing one of the following: 1) Selecting one or more disks and clicking the “Add” button; 2) Clicking the “Add All” button; 3) Double-clicking on a disk in the “Available disks” list; 4) Dragging-and-dropping one or more disks from the “Available disks” list to the “Selected disks” list.

“OK” button—This button becomes enabled when the user selects at least one disk from the “Available disks” list. Validations are performed when the button is clicked to ensure that: the Name field is unique; the Label field is unique; the volume size is less than the total free space for all the selected disks; and the volume size is large enough to hold all the mailbox database files assigned to the volume.

FIG. 13 illustrates a panel that is displayed in Custom mode to allow a user to select one or more disks to set aside as hot spare. The components function similarly as described above with reference to the “Assign disks for Volume” components of FIG. 12. The disks shown in the “Available disks” list that are selectable are empty. Once the selected disks have been marked as hot spare, they cannot be assigned to any volume.

FIG. 14 illustrates a panel that is displayed in Custom Mode to allow a user to edit an EDB or S™ mailbox database file. The various components shown on the panel are described as follows:

“Mailbox database name”—Database name for both EDB and STM.



## 11

“Size field”—Size of the EDB or S™ file.

“User profile”—The user profile type.

“Database file name”—Name of the mailbox database file for both EDB and STM.

“OK” button—When the “OK” button is clicked, validations are performed to ensure that: the mailbox database name is unique and non-empty; the database file name is unique and non-empty; and the file size is not larger than the volume capacity

FIG. 15 illustrates a panel that is displayed in Custom Mode to allow a user to edit a Transaction Logs entry. The various components shown on the panel are described as follows:

“Transaction Logs”—A read-only field that is automatically generated by concatenating the name of the Storage Group and “TransLogs” constant string which is defined in a properties file.

“Transaction Logs size”—Size of transaction logs.

“User profile”—The user profile type.

“OK” button—When the “OK” button is clicked, a validation is performed to ensure that the transaction logs size is not larger than the volume capacity.

FIG. 16 illustrates a panel that is displayed in Typical Mode after the user has completed configuration for each Storage Group. The panel allows the user to select disks for creating the Storage Groups. The user can select one or more disks from the “Available disks” list to the “Selected disks” list. The components function similarly as described above with reference to the “Assign disks for Volume” components of FIG. 12.

FIG. 17 illustrates a panel that is displayed in both Typical Mode and Custom Mode. The panel allows the user to select empty volumes to use. In Typical Mode, the panel of FIG. 17 is displayed after the panel of FIG. 16, if there is at least one disk that contains empty volumes (not assigned to any existing Storage Groups). In Custom Mode, the panel of FIG. 17 is displayed after the panel of FIG. 7, if at least one empty volume is found.

FIGS. 18 and 19 illustrate panels that display different views of the storage configuration or layout that results after stepping through the above-described panels in Typical Mode or Custom Mode. FIG. 18 illustrates a Storage Group view of the layout. (As shown in FIG. 18, the “Storage Group View” tab at the top of the panel is currently selected.) The Storage Group view is based on existing and new Storage Groups. The Storage Group view lists all mailbox database files that will reside on new volumes.

FIG. 19 illustrates a Volume view of the layout. (As shown in FIG. 19, the “Volume View” tab at the top of the panel is currently selected.) The Volume View displays a tree view starting from new disk groups that will be created during implementation of the layout. In each disk group, new volumes belonging to the disk group are displayed. Information about disks and mailbox database files is also displayed.

FIG. 20 illustrates a panel that is displayed upon completion of the Storage Configuration wizard. The user can select options to implement the suggested layout and/or save the suggested layout as an XML file. In response to clicking the “Finish” button, the selected options are performed, i.e., the layout is implemented and/or saved. If the user chooses to save the suggested layout as an XML file but not implement the layout, the user may be able to open the layout at a later time and implement the layout at the click of a button (possibly after modifying the layout as desired).

The panel indicates actions to be performed when the layout is implemented. These actions include:

## 12

Moving Storage Groups—The panel provides a listing of all the existing Storage Group that will be moved from Basic Disk Group to Dynamic Disk Group.

Delete empty volumes—The panel provides a listing of all empty volumes that will be deleted during implementation of the layout.

Create new Disk Groups—The panel provides a listing of all Disk Groups that will be created during implementation of the layout.

Create new volumes—The panel provides a listing of all volumes that will be created during implementation of the layout.

FIG. 21—Configuration Software Architecture

FIG. 21 illustrates a software architecture of exemplary network configuration software according to one embodiment. The network configuration software may execute on a client computer system 102 used by a network administrator. The network configuration software may provide a management console to the network administrator for performing various network configuration and/or management functions. In one embodiment, the network configuration software shown in FIG. 21 includes one or more software programs (referred to collectively as the “Storage Configuration” program) for automatically determining storage configurations or storage layouts for computers in the system. The network configuration software shown in FIG. 21 may also include software for automatically applying a determined storage configuration to one or more computers in the system. In the currently preferred embodiment, the network configuration software of FIG. 21 is offered by VERITAS Software Corporation (VERITAS Software).

As shown, the software architecture in FIG. 21 may include an Administrator GUI or console 152. The Administrator GUI 152 may support various extensions, such as a VM (Volume Manager) Extension 154 and an Edition Extension 156. The VM Extension 154 provides a console for volume management for non-volatile storage, e.g., disk drives. The Edition Extension 156 provides the user interface for various Edition functionality, such as the Storage Configurator. An Edition Installer 158 may be provided for installing the Edition Extension 156. An Edition Extension (EE) Backup application 160 may also be provided for backing up the application’s data. A SCP Interface 150 may also be provided for providing external access to Storage Configurator Provider 178.

The Administrator GUI 162 may interface to a Client Infrastructure layer 172. The Client Infrastructure layer 172 may in turn interface to a Server (Bus) program layer 174. The Server (Bus) program layer 164 may abstract various communication functions among various distributed software programs executing in the network system. As shown, the Server (Bus) program layer 164 may interface to various programs, such as VM (Volume Manager) Provider 172, VAIL (Veritas Array Integration Layer) Providers 174, and an Exchange Provider 176. The VM provider 172 provides volume management for non-volatile storage, e.g., disk drives. The VAIL Providers 174 provide management services for RAID arrays. The Exchange Provider 176 provides management services for Microsoft Exchange.

The Administrator GUI 162 may also interface to a SC (Storage Configurator) Provider 178. The SC Provider 178 may include at least a portion of the Storage Configuration software of one embodiment of the invention. The SC Provider 178 may interface to an application specific layer 180 used in automatically determining a storage configura-



tion for a specific application **182**. The application specific layer **180** may in turn interface to a specific application **182**. The application specific layer **180** preferably includes functionality that is specific to the application **182** to which the layer **180** corresponds.

In the embodiment shown, the application **182** is Microsoft Exchange, and the application specific layer **180** is the Exchange Edition Layer (EEL). Due to COM interface requirements of Microsoft Exchange, the EEL **180** interfaces to Microsoft Exchange through a COM object layer referred to as CDOEXM (Collaborative Data Object Exchange Management Layer) **181**. The CDO Exchange Management Layer **181** is an administrative layer that uses the Collaborative Data Object API from Microsoft to obtain information and configure Microsoft Exchange.

In one embodiment, the SC Provider **178** includes core functionality that is common or generic to any application, and each application specific layer **180** includes application specific functionality for its specific application. This architecture enables modular reuse of the SC Provider **178** for different applications. This architecture obviates the need for a developer to write a new optimization engine for each application in which a storage configurator is desired. Thus the SC Provider **178** operates to abstract the general storage layout optimization algorithms from the application specific rules which are contained in the application specific layer **180**.

FIG. 22—Storage Configurator Software Architecture

FIG. 22 illustrates a software architecture of the Storage Configurator software according to one embodiment of the invention. As shown, the Storage Configurator software includes a Storage Configurator Provider **178** and various other files. The Storage Configurator provider **178** includes a Configurator Core **202** and an XML parser **204**. The XML parser **204** performs XML parsing functions. The Configurator Core **202** may perform various operations to obtain information used in determining an “optimal” storage layout for the system.

The Storage Configurator software also includes a FreeSpace.xml file **222**, a Priority.xml file **224**, and a Layout.xml file **234**.

The FreeSpace.xml file **222** contains information about the available free space in the storage system (or network system). In one embodiment, the FreeSpace.xml file **222** contains information about all of the free space in the system. The FreeSpace.xml file **222** is created by a Discoverer component **212**, as discussed below.

The Priority.xml file **224** may also be referred to as a “rule information” data structure or “rule data structure”. The Priority.xml file **224** comprises rule information used in determining the storage configuration. The rule information may comprise rules about storage configuration principles of the application. The rule information may also comprise priority information, wherein the priority information specifies a priority in application of rules in the rule information. In one embodiment, the rule information comprises storage types in order of priority for different application data types. The Priority.xml file **224** (or rule information data structure) may specify relative preferences for each type of storage to each type of application data type or volume. The priority.xml document thus stores the priority or rule information for the layout determination. In a Microsoft Exchange example, the Priority.xml file **224** specifies relative preferences for each type of storage to each type of Exchange

volume. In one embodiment, the Priority.xml file **224** is not modifiable by the user. In another embodiment, the Priority.xml file **224** is modifiable by the user. In another embodiment, a plurality of Priority.xml files **224** are maintained in the system, and the user may select a preferred Priority.xml file **224**. Alternatively, the type of application for which storage configuration is being determined may specify one of (or a subset of) the plurality of Priority.xml files **224**.

Exemplary rules are listed below:

- Each Storage Group results in its own Disk Group
- Each class of user results in its own Database
- Transaction Logs kept on separate set of spindles
- All STM files within Storage Group can be grouped together
- All EDB files within Storage Group can be grouped together
- Last resort: Put STM and EDB files from same Storage Group together
- Suggested STM file size=# users\*mailbox size
- Suggested EDB file size=# users\*mailbox size
- Additional rules about hot spares, RAID-5 logging, FMR logs, etc.

These rules may be used in automatic creation of a storage configuration (e.g., Typical mode of the Storage Configurator program). However, the user may also override these rules if desired (e.g., Custom mode of the Storage Configurator program).

As described above, the user of the storage configurator provider **178** may provide input to a GUI about the target application, including various information regarding how the application will be used. A UserInput.xml file **232** may be created in response to the user input supplied by the user.

The user input (UserInput.xml file) **232** is received by the Application Specific Layer **180**, in this example the Exchange Layer or EEL. The Application Specific Layer **180** contains knowledge about the specifics of its respective application. In a current example where the application is Microsoft Exchange, the EEL **180** contains knowledge of the specifics of Microsoft Exchange and operates to create a logical exchange storage configuration based on inputs from the user. The logical exchange storage configuration may contain the name and size of each exchange object requiring storage. The Application Specific Layer **180** generates the Layout.xml document **234**. In this example, where the application is Microsoft Exchange, and the Layer **180** is the EEL, the document is referred to as Exchange Layout.xml **234**.

In one embodiment, the EEL may automatically decide how much space to allocate each file. In another embodiment, the user may specify desired file sizes. The EEL may receive a user input data structure from the GUI and convert this information into the application information. The EEL includes various application specific rules.

As one example for Microsoft Exchange, the EEL may comprise the following rules:

- Possibly split databases if they are too big;
- Number of parallel backups used as minimum for number of Storage Groups
- If content indexing turned on, add 20% to above numbers

Transaction Log size=(Total size of all databases within Storage Group\*0.20)+100 MB

Memory constraints:

RAM>=(130 MB\*SGCount)+((DBCCount-SGCount)\*25 MB)+126 MB



## 15

If total mailboxes <500, RAM >=256 MB  
 If 500<=total mailboxes <1000, RAM>=512 MB  
 If 1000<=total mailboxes <2500, RAM>=1 GB  
 If total mailboxes>=2500, RAM>=2 GB

The Configurator Core **202** includes a Discoverer **212**, a Core Logic engine **214**, and an Implementer **216**.

The Discoverer **212** operates to automatically detect existing storage in the network system. The Discoverer **212** generates the FreeSpace.xml file **222**, described above. The FreeSpace.xml document **222** describes the free space or available storage in the storage system.

The FreeSpace.xml document **222** is available to the Core Logic engine **214**. The Core Logic engine **214** receives the free space information from FreeSpace.xml **222**. As noted above, the Storage Configurator software also includes a priority document referred to as Priority.xml **224**. The optimization algorithm(s) in the Core Logic Engine **214** analyze the priority or rule information in Priority.xml **224** in determining how to optimize usage of the discovered free space or available storage space in the system.

The Core Logic Engine **214** analyzes the available free space in the FreeSpace.xml document **222**, analyzes the user input embodied in the Layout.xml document **234**, and examines the priority or rules contained in the Priority.xml document **224**, and operates to generate a suggested layout or suggested storage configuration for the existing storage. The suggested layout is embodied in a document called SuggestedLayoutFile.xml **226**. The Core Logic Engine **214** preferably forms the core of the storage configurator. In one embodiment, the Core Logic Engine **214** may utilize a policy whereby the volume size is equal to the sum of the sizes of all files to be placed on the volume.

The Implementer **216** receives the SuggestedLayoutFile.xml **226** and operates to implement the suggested layout. In other words, the Implementer **216** analyzes the information in SuggestedLayoutFile.xml **226** and configures the various network resources (including non-volatile memory or disks) according to the layout information, optionally in response to user input or approval of the suggested layout.

Therefore, in determining a (preferably optimal) storage configuration according to one embodiment, the Core Logic Engine **214** utilizes user input received from the user, such as through a console wizard or GUI; the amount of available storage or free space (or memory, e.g., non-volatile memory and/or RAM) present in the storage system (e.g., one or more servers); and the Priority.xml file.

FIG. 23—Flowchart Diagram

FIG. 23 is a flowchart diagram illustrating one embodiment of a method of operation of the Storage Configurator software to automatically determine an optimal storage configuration for an application (e.g., the Typical mode of the Storage Configurator).

In **302** the method may first receive information on existing storage capabilities in the storage system. For example, the user may enter information in a graphical user interface (GUI) regarding the existing storage capabilities in the storage system. Alternatively, the method may automatically (i.e., programmatically) determine or detect existing storage capabilities in the storage system. Step **302** may be performed by the Discoverer program **212** of FIG. 22. In an alternate embodiment, the user may manually provide at least a portion of the information regarding the available or

## 16

existing storage in the system. A flowchart diagram illustrating operation of the Discoverer **212** is provided in FIG. 25.

In **304** the method may receive information regarding the target application. For example, the user may enter information in a graphical user interface (GUI) regarding the application. The application layer **180** of the storage configurator may also, or instead, query the application to obtain information regarding the application. The information regarding the application may comprise information regarding various application objects (e.g., files, databases, etc.) present in the application. As noted above, the storage configurator software program may present a graphical user interface, such as a GUI of a console wizard, to the user. The user may then input values which will be used to aid in determining the optimal storage layout. Table 1 below describes exemplary questions that may be presented by the GUI in **304**.

TABLE 1

Question	Influences
Content Indexing	Size of the volume increases
Number of groups of users (classes of users)	Total number of Mailbox Databases
Number of users in group * size of each mailbox	Size of each Database
Maximum size of Backup	Maximum size of Database
Number of parallel Backups	Minimum number of Storage Groups
Which LUNs to include	Maximum number and size of Databases we can configure (and where to put databases and transaction logs). Also influences which empty volumes can be included.
Which Volumes to include	Where we can configure databases and transaction logs. Any included volumes which are used will be deleted and their space will be reallocated.

Alternatively, information regarding the application may be obtained from a previously created data structure. For example, the user (or another user) may have previously created an application data structure that contains the application information. In another embodiment, application program vendors may provide such an application data structure with their applications.

In **306** the method may automatically determine a storage configuration (or storage layout) for the storage system based on the existing storage capabilities in the storage system and the information regarding the application. In automatically determining the storage configuration, the method may examine a first data structure (called the “rule data structure” or “priority.xml”) comprising rule information. The rule information may comprise rules about storage configuration principles of the application.

In one embodiment, the rule information comprises priority information that specifies a priority in application of the rules. For example, the rule information may comprise a structure or list of storage types in order of priority for different application data types.

In one embodiment, the method automatically determines the storage configuration by examining available storage in the storage system, determining storage allocations based on the available storage in the storage system and the rule information, and generating the storage configuration based on the storage allocations.

More specifically, one embodiment may operate as follows. The method may automatically determine the storage configuration by first performing pattern matching between



the application information and the rule information. This pattern matching may comprise matching the contents of the rule information data structure with the contents of the application information. The method may then create a second data structure (“skeleton” data structure) comprising storage organization information of the storage configuration based on the results of the pattern matching. The second data structure may comprise information regarding which application objects should be in which storage subsets (e.g., disk groups). The method may further determine groupings of application objects. The method may then automatically determine a storage configuration using two or more of the first data structure, the second data structure, and the groupings of application objects.

The method may operate to determine the optimal storage configuration (or storage layout) based on the existing storage in the system determined in 302, the information received from the user in 304, and the priority.xml document referred to above. In one embodiment, step 306 is performed by the core logic engine 214 described in FIG. 22.

In step 306 the operation of determining a storage layout is performed automatically. In the present application the term “automatically” refers to a function that is performed programmatically, i.e., the function is performed primarily by software executing on a computer system. Thus, in one embodiment the user is not manually required to specify the storage layout. Rather, based on some simple user input and other information, in step 306 the storage configurator program applies an allocation algorithm (e.g., various optimization rules) to automatically or programmatically determine the optimal storage layout based on this information.

In 308 the method may optionally operate to display the determined storage configuration to the user. The display may also include a graphical user interface element, such as a dialog box or other element where the user may graphically chose to accept or decline the determined optimal storage layout.

As shown in the flowchart of FIG. 23, if the user does not accept the determined optimal storage layout, then operation of the method completes. In this case, the user may then repeat the method above or may chose to manually configure a storage layout. If the user accepts the optimal storage layout in step 310 then in step 312 the implementer software 216 may operate to automatically or programmatically implement the layout. This may include placing the target application’s data objects.

In one embodiment, the method may simply store the determined storage configuration in memory, possibly without displaying the determined storage configuration. In yet another embodiment, the method may automatically apply the determined storage configuration, again possibly without displaying the determined storage configuration or receiving any acceptance from the user.

FIG. 24—Alternate Embodiment

FIG. 24 illustrates an alternate embodiment of the method described in the flowchart of FIG. 23. In the flowchart of FIG. 24, in step 303 the user may also choose to specify one or more constraints on the storage configuration or storage layout prior to the optimal storage configuration being determined in step 306. These constraints are received by the core logic engine 214 and are used in step 306A to determine the (preferably optimal) storage layout. In other words, in step 306A, the optimal storage layout is determined based on information described above (including the available amount of storage or free space, user input regarding the

target application, and the priority.xml document) and is also determined based on the constraints specified by the user. Thus, the determined optimal storage layout produced in step 306 will be determined based on an accounting for the constraints identified by the user. This provides a mechanism for the user to specify constraints on priority and to have a storage layout automatically determined taking these constraints into account.

In each of the embodiments of FIGS. 23 and 24, the user may choose to specify some of the input provided to steps 306 and/or 306A and repeat the above process to generate new optimal storage layouts based on this new information. For example, the user may re-invoke the wizard GUI and provide new information about the target application in step 304, and then rerun step 306 to determine a new optimal storage layout based on the new information provided by the user in step 304. As another example, if the user adds more storage to the system, the user may repeat the above method, whereby the optimal storage layout determined in step 306 may differ due to the fact that a greater or lesser amount of storage was added to and/or removed from the system in 302. As yet another example, the user may select a new priority.xml document, or modify an existing priority.xml document, and rerun the method of FIG. 23 to determine a new optimal storage layout based on this newer modified priority.xml document. In some embodiments, the storage configurator program may be automatically re-invoked in response to certain events, such as the addition or removal of storage devices from the storage system.

#### Modification of Optimal Storage Layout

The optimal storage layout displayed in step 306 may be displayed in any of various ways. In one embodiment, the optimal storage layout is displayed as a tree structure. In another embodiment, the optimal storage layout may be shown iconically in a diagram. For example, a storage configuration diagram may be displayed which includes icons for each of the storage units, e.g., server units, SANs etc. and the optimal storage layout may be graphically and/or iconically depicted in this storage configuration diagram. Regardless of the manner in which the optimal storage layout is displayed to the user, in one embodiment the user may graphically manipulate the displayed storage layout to make modifications to the storage layout. Thus, if the optimal storage layout is displayed as a tree structure in step 306, the user may optionally provide the ability to manually adjust different aspects of the tree structure or different values associated with the tree structure to change the storage layout. For example, the user may make changes and then hit “apply” to apply these changes. If the optimal storage layout is displayed graphically or iconically, the user may use various drag and drop techniques or other graphical editing techniques to change the optimal storage layout as desired. In another embodiment, the optimal storage layout determined by the storage configurator 178 is not modifiable by the user.

FIG. 25—Discoverer Operation

FIG. 25 is a flowchart diagram illustrating operation of the discoverer component 212 in the storage configurator provider 178. The flowchart diagram of FIG. 25 also illustrates operation of step 302 of FIGS. 23 and 24.

As shown, in 402 the method queries the disk objects registered by the volume manager. This operates to determine what disk objects are present in the system.

In 404 the method determines if each disk object determined in 402 is part of a hardware array. A portion of this



operation may be performed by VAIL Provider 174. If a disk object is determined to be part of a hardware array, then in 406 the method obtains the disk object's respective hardware array properties. Operation then advances to step 408. If a respective disk object is determined to not be part of a hardware array in step 404, then operation proceeds directly to 408.

As shown, in 410 and 412 the method queries each of the disk objects about its properties. The property information for the disk objects is then converted into a FreeSpace.xml document in step 414. The FreeSpace.xml document contains information about free space in the system, preferably all free space available to the application within the system.

FIG. 26—Automatically Determining a Storage Configuration

FIG. 26 is a flowchart diagram illustrating automatically determining a storage configuration according to one embodiment. FIG. 26 provides more detail regarding steps 306 and 306A of FIGS. 23 and 24, respectively.

In 442 the method may first perform pattern matching between the application information and the rule information. This pattern matching may comprise matching the contents of the rule information data structure with the contents of the application information. The method may then create a second data structure (also called the "skeleton" data structure and the "NewTree" data structure) comprising storage organization information of the storage configuration based on the results of the pattern matching. The second data structure may comprise information regarding which application objects should be in which storage subsets (e.g., disk groups). The second data structure may be an intermediate form of the SuggestedLayout data structure comprising storage organization information. In one embodiment, each of the rule information data structure (priority.xml) 224 and the application information 234 may comprise text files. The method may perform pattern matching or textual analysis between these two text files to determine a matching or correspondence between data in each of the two text files.

In 444 the method may further determine groupings of application objects. This grouping information may affect volume sizes to be allocated. The method may determine respective groupings of application objects in 444, and then in 522 the method may examine the available storage and apply an allocation algorithm based on the grouping. Steps 444 and 522 may be performed two or more times in a repetitive or iterative fashion, as shown. In other words, the method may determine a first grouping of application objects and apply the allocation algorithm to this first grouping, and then determine a second grouping of application objects and apply the allocation algorithm, and so on, until the optimal (or an acceptable) storage configuration is produced. Steps 444 and 522 are described further with respect to steps 504-512 and 522A-E of FIG. 27. The allocation algorithm may be either a "top down" or "bottom up" greedy (or pseudo-greedy) algorithm with backtracking. The results of these steps may produce a storage configuration. Thus, in one embodiment, the storage configuration may be automatically determined using at least a subset of the above information, e.g., two or more of the first data structure (rule data structure), the second or skeleton data structure, and the groupings of application objects.

After the storage configuration has been determined or generated, in 448 the storage configuration may be stored in memory. The storage configuration may be useable in con-

figuring the storage system. The method may then operate to configure the storage system according to the storage configuration.

FIG. 27—Core Logic Engine

FIG. 27 is a flowchart diagram illustrating one embodiment of operation of the core logic engine 214. In FIG. 27, step 502 corresponds to step 442 of FIG. 26, steps 504-510 correspond to step 444 of FIG. 26, and steps 522A-E correspond to step 522 of FIG. 26.

As used herein, the term application object type refers to items in the application which require disk space. Examples include transaction logs, EDBs and STMs. In an exemplary embodiment where the application is Microsoft Exchange, exemplary application object types, (in this case referred to as exchange object types) would include transaction logs, Exchange DBs, and STMs.

As used herein, the term "dedicated spindle set" refers to disks or non-volatile storage elements which contain data for a given set of application object types and store no other significant data.

In step 502 the method examines the priority.xml document and the application specific layer 180 (e.g., the Exchange layer 180). In step 502 the method examines this information and performs pattern matching or text string matching in order to create a data structure or skeleton of the constraints of the layout to be created (the second data structure). In step 502 the method compares the text string entries in the priority document with the text string values created in the Exchange layout.xml file 234 to identify matches between these two documents. In response to these matches, the method then creates a second data structure, also referred to as a skeleton data structure, with the beginning organization of the storage configuration or layout. An example of operation of step 502 is discussed below.

In steps 504-510 and 522A-E the method determines groupings of application objects and attempts to allocate storage to these groupings to create the storage configuration. The allocation algorithm performed in 522 (522A-E) may be a "greedy" algorithm with back tracking, although other types of allocation algorithms may be performed. The allocation algorithm may be a top down algorithm or a bottom up algorithm, or a combination thereof. Any of various types of allocation algorithms may be used, as desired.

As shown, in step 504 the method uses the freespace.xml information to attempt to find a dedicated spindle set for each application object type. In performing step 504, the method performs an allocation algorithm in step 522 (referred to here as 522A) to attempt to find disk space for the layout. If the algorithm in step 522A locates disk space sufficient for a dedicated spindle set for each application object type, then the search completes in step 232.

If sufficient disk space is not found in 522A, then in step 506 the method attempts to find a dedicated spindle set for transaction logs and a dedicated spindle set for each set of application object types. Again, in performing step 506, the method performs the allocation algorithm of 522 (referred to here as 522B). If space is found, then the search completes in 232.

If sufficient space is not found in step 522B, then operation advances to step 508. In step 508 the method attempts to locate a dedicated spindle set for transaction logs, and a single dedicated spindle set for all application object types.



## 21

Again, the algorithm of **522** (referred to here as **522C**) is performed. If space is found, then the search completes in **232**.

If sufficient space is not found in step **522C**, then operation advances to step **510**. In step **510** the method attempts to locate a dedicated spindle set for transaction logs, and a single spindle set for all application object types. In step **510**, the method considers free space on disks with existing storage volumes (as opposed to step **508**, which only considers disks without existing storage volumes). Again, the algorithm of step **522** (referred to here as **522D**) is used in this process.

If sufficient space is not found in step **522D**, then the method proceeds to step **512** to attempt to locate non-fault tolerant layouts. In step **512** the method attempts to determine any non-fault tolerant layouts using the algorithm of step **522** (referred to here as **522E**). If a fault tolerant layout is determined in step **512**, then the search ends in **232**. If no non-fault tolerant layouts are able to be determined in step **512**, then the configuration fails as illustrated in **234**.

In an alternate embodiment, the methodology described in FIG. **27** is reversed, and a “bottom up” algorithm is used. In other words, all non-fault tolerant layouts are first examined to determine possible layouts, then operation proceeds to progressively improve storage layouts through step **510**, then performs steps **508**, **506** and **504**. In this embodiment, once one of the operations fails to produce a valid result, then further operations need not be performed. In other words, if it is not possible to find a dedicated spindle set for transaction logs, in a single dedicated spindle set for all application object types and STMs together in step **508**, then there is no point in performing either steps **506** or **504**, since these solutions require even greater available disk free space than step **508**.

In one embodiment, the allocation algorithm uses a combination of bottom up and top down approaches. For example, the method may begin the analysis from a most expensive to least expensive file grouping (as shown in steps **504-510**). For each determined file grouping (in one of steps **504-510**), the method performs the allocation algorithm by starting with the lowest priority and least expensive volume option in priority.xml file and then working upward. As noted above, the allocation algorithm is performed to actually allocate volumes based on the file grouping. Thus the method may proceed through file groupings in a top down manner, and the allocation algorithm may proceed in a bottom up manner. After the method has found a volume for the sum of all file groupings, then an improvement pass may be performed, as described further below.

The following describes a currently preferred embodiment of a portion of the storage configuration creation process (i.e., steps **504-510** and **522A-E**):

1. Obtain an initial grouping of objects, e.g. how files can be grouped on volumes. Specifically, the first pass generally attempts to keep each file on its own volume (and, therefore, it’s own spindle set) to help optimize performance.
2. The allocation algorithm attempts to “create” and then map suggested volumes to the objects. “Create” in this sense does not mean to actually create the physical volume. Instead, the term “create” in this context means to reserve space from the “freeSpace.xml” data structure for the chosen volume.
  - a. In the first pass, the storage configurator chooses the least expensive (least resource intensive) volume type available that is still fault tolerant—often a

## 22

simple concatenated volume over hardware RAID, or a simple software mirror over a JBOD.

- b. The allocation algorithm attempts to create and assign a volume to each group of files:
  - i. If the allocation algorithm runs out of space (e.g., all the available elements in freeSpace have been marked as reserved and there are still files that need volumes), the allocation algorithm returns to step 1 and obtains a different, less expensive, file grouping. It is generally easier to create a few big volumes rather than many smaller, specialized volumes. The very last time, the allocation algorithm will allow the creation of non-fault tolerant volumes as well.
  - ii. If the allocation algorithm succeeds in finding an inexpensive volume for every file grouping, the allocation algorithm performs an “improvement pass”. This pass goes through one disk group at a time, releases all the reserved objects for that disk group from FreeSpace, and then it tries to create a volume suggestion starting with the most expensive volume types, and working its way to the least expensive. This way, the allocation algorithm will produce at least as good of a layout as the starting layout, but the layout has a good likelihood of being improved.

## FIGS. 28-31

FIG. **28** is a flowchart diagram illustrating more detailed operation of the Core Logic Engine **214** as described in FIG. **27** above. As shown, in **602** the DoSuggestLayout( ) function is called. The DoSuggestLayout( ) function is essentially the “main( )” of the Core Logic Engine **214**. The Core Logic Engine **214** receives the FreeSpace.xml file, i.e., the structure which contains all the information about the disks that are available for use by the application’s storage needs, and an application specific Layout.xml file that contains the object hierarchy and space requirements for the application objects that need storage. The third input, the Priority.xml file, contains template information about which volume types are best for each application object type. The Priority.xml file is preferably a static file and is saved to disk at install time. The Priority.xml file is read from disk on the fly in MakeNewTree( ) so that the storage configurator can “understand” the rules for allocating space for the application’s needs.

As shown, the DoSuggestLayout( ) function calls the MakeNewTree( ) step **502** (step **502** of FIG. **27**). The MakeNewTree( ) function creates a new XML DOM tree by intelligently recognizing patterns in layout and priority. The MakeNewTree( ) function is described further in FIG. **29**.

The DoSuggestLayout( ) function also calls the CreateLayout( ) function **606**. The CreateLayout( ) function **606** determines an optimal layout for the objects specified in “newTree”. The CreateLayout( ) function **606** corresponds to steps **504-512**, and **522A-E** of FIG. **27**.

The DoSuggestLayout( ) function also calls the GenerateSLOutputFile( ) function **608**. This function formats the layout for consumption as specified in SuggestedLayout.xsd.

FIG. **29** is a flowchart diagram illustrating more detailed operation of step **502**. The operation in **502** creates a new XMLDOM tree by intelligently recognizing patterns in layout and priority. As shown, in **622** the method first verifies form of the inputs, ApplicationLayout.xml **234** and Priority.xml **224**. In **624** the method finds the “key” element



of ApplicationLayout.xml 234. In 626 the method finds all elements with an attribute that matches this key from both the Priority.xml and Application.xml data structures 224 and 234. In 628 the application specific objects that require disk space are first divided into disk groups as specified in the Priority.xml file 224. This is performed since the largest division of storage is on the disk group level. In 630 the method adds objects to the “newTree” data structure (“second data structure” or “skeleton data structure”).

It is noted that different priority and layout files will yield different results. In the operation of this function, the “Key” elements of the two files are preferably equal. Both files should have elements with this “Key” attribute and these elements should have the same names in each file. This is how the storage configurator determines the structure for which it will be creating volumes. The priority file preferably has a “DiskGroup” element as well. This is how the storage configurator determines how to lay out Disk groups and volumes within the disk group. The newTree data structure is something of a combination of layout and priority that the storage configurator uses internally. The Priority.xml file describes how Exchange volume types and disk groups fit together. The ExchangeLayout.xml file represents the requested configuration in terms of Exchange. Since disk groups by definition need to have their own spindle set, the storage configurator preferably understands how disk groups fit into the Exchange configuration. The MakeNewTree( ) function takes these two files and merges their structures into newTree to reflect how disk groups fit into the Exchange picture.

FIG. 30 is a flowchart diagram illustrating operation of the CreateLayout( ) function. The CreateLayout( ) function performs the actual generation of the layout of volumes to disks and application files to volumes.

As shown, in 642 the method finds the disk group element in the NewTree( ) data structure. It is noted that disk space will be allocated one disk group at a time. Steps 644 onward essentially perform steps 504, 506, 508 510, 512 and 522A-E of FIG. 27.

In 644 the method finds node types that can be shared from Priority.xml. In 646 the method obtains objects that need new space from the NewTree data structure.

In 444 the method creates a file grouping, as discussed above with respect to step 444 of FIG. 26.

In 522 the method executes the AllocateSpace( ) function to implement the allocation algorithm.

In one embodiment of steps 444 and 522, the method attempts to find a dedicated spindle set for each application object type (FIG. 27, 504). This is performed by getting lists of objects, where each list contains the set of objects that will reside together on a volume. Each list is a subset of “otNodes” (object type nodes). Then, the method calls “AllocateSpace( )” for each list of objects. The “AllocateSpace( )” method is the allocation algorithm performed in 522 (522A-522E) This process continues until either disk space has been found for every object in the otNodes list, or there is no more disk space available. Note that this algorithm preferably never creates more than one volume on a single set of disks.

After steps 444 and 522, i.e., after a volume has been found for all disk groups, in 650 the method may perform and “improvement pass” per disk group. As described above, if the allocation algorithm succeeds in finding an inexpensive volume for every file grouping, the allocation algorithm preferably performs an “improvement pass”. This pass examines one disk group at a time, releases all the reserved objects for that disk group from FreeSpace, and

then attempts to create a volume suggestion starting with the most expensive volume types, and working its way to the least expensive. This way, the allocation algorithm will produce at least as good of a layout as the starting layout, but the layout has a good likelihood of being improved.

The improvement pass performs operations similar to those described in step 672 below, except the improvement pass starts from top to bottom, instead of bottom to top.

FIG. 31 is a flowchart diagram illustrating operation of the AllocateSpace( ) method. This preferred embodiment of the AllocateSpace( ) method performs a combination of a bottom up and top down algorithm with a refinement layer to find disk space. The method in FIG. 31 is preferably performed for each file group and/or volume.

As shown, in 662 the method examines the type of objects in the list of files from the current file group. In 664 the method gets the sum of the size of the data objects that will be placed on the volume. In 666 the method marks all objects to be placed on volume as “held” in NewTree. In 668 the method finds the name of the disk group to which this volume will belong. In 670 the method finds the corresponding object type value in the Priority.xml data structure.

In 672 the method selects disks for the volume as described below. As described above with respect to FIG. 27, the allocation algorithm attempts to “create” and then map suggested volumes to the objects. “Create” in this sense does not mean to actually create the physical volume. Instead, the term “create” in this context means to reserve space from the “freeSpace.xml” data structure for the chosen volume. In the first pass, the storage configurator chooses the least expensive (least resource intensive) volume type available that is still fault tolerant—often a simple concatenated volume over hardware RAID, or a simple software mirror over a JBOD.” This is performed by traversing the disk specifications for the given file type in the Priority.xml file from the bottom up. Each volume type specified in the Priority.xml file contains both a hardware RAID type and a software RAID type specification. First, the algorithm finds all of the available disk objects in FreeSpace.xml that are of the given hardware RAID type.

If there are disks of the specified hardware RAID type, the algorithm checks the following:

a) if there are enough spindles of that type available to create the software RAID volume (i.e., at least 2 spindles are required for software mirrored or striped volumes, and at least 4 spindles are required for software striped mirrored volumes).

b) if there are enough spindles available, the algorithm proceeds to calculate the actual amount of disk space needed given the requested volume size for the data, and the extra requirements a given software RAID volume might need. For example, if the specified software volume type is a mirrored volume, by default the storage configurator will enable dirty region logging for more optimal volume performance. In this case, the algorithm will add the amount needed for the log to the total volume size.

c) then the algorithm passes the list of LUNs (a.k.a. disks or spindles), the total amount of space needed, and the type of software RAID specified to a PickLUNs( ) function. The PickLUNs( ) function finds a subset of those LUNs that fulfills the volume requirements, and then reserves those LUNs for the given volume. The PickLUNs( ) function thus finds a set of disks that will fulfill the volume requirements.



FIG. 32—Implementor

FIG. 32 is a flowchart diagram illustrating operation of the implementor 216. As shown, in step 702 the implementor 216 parses the suggested layout.xml file and obtains a list of disk groups, volumes and hot spares to be created.

In step 704 the method determines if all required disks are still available. If all required disks are determined to not be still available in step 704, then in step 706 an error is recorded for the report to be generated. Operation then proceeds to step 716 where the error report is returned. In this instance, the error report will indicate an error has occurred, since all required disks are not available to implement the layout in the suggested layout.xml file.

If all required disks are determined to still be available in step 704, then in step 708 the method resolves any duplicate names of volumes. In step 710 the method creates disk groups. In step 712 the method creates volumes and directories on each volume. In step 714 the method creates any hot spares. In step 716 the method returns an error report. If the implementor 216 successfully implemented the suggested layout, then the error report will return indicating no error.

Example Data Structures

Exemplary FreeSpace.xml

Returned by the Discoverer component;  
Contains all free space on the system.

```
<FreeSpaceDescFile>
  <!--FreeSpaceDesc file generated by Discoverer-->
  . . . HEADER . . .
  <StorageLayout>
    <LocallyAttachedStorage>
      <LUN>
        <ID>
          <DisplayName>Harddisk7</DisplayName>
          <InternalID><![CDATA[SCSI\DISK&VEN_
            IBM&PROD_DDRS-
            34560D&REV_DC1B\5
          </ID>
          <TotalSpace>4565030400</TotalSpace>
          <FreeSpace>4565030400</FreeSpace>
          <IsDynamic>>false</IsDynamic>
          <IsBCV>>false</IsBCV>
          <IsFaultTolerant>>false</IsFaultTolerant>
          <TypeOfRAID>NORAI</TypeOfRAID>
        </LUN>
```

Exemplary UserInput.xml

Output by a GUI;  
Contains all values the user typed in

```
<?xml version="1.0"?>
<UserInputs xmlns:xsi="http://www.w3.org/2000/10/XMLSchema-
instance" xsi:
. . . HEADER . . .
  <SpecialFlags>
    <GetExistingExchangeInfo>>true</GetExistingExchangeInfo>
    <Caller>RuntimeUI</Caller>
    <CreateStorageGroup>>false</CreateStorageGroup>
    <CreateDatabase>>false</CreateDatabase>
    <StorageGroupName/>
  </SpecialFlags>
  <UserPriority>MaxCapacity</UserPriority>
  <ClusteringEnabled>>false</ClusteringEnabled>
  <UserGroup>
    <Name>TechSupport</Name>
    <NumberOfUsers>30</NumberOfUsers>
```

-continued

Exemplary UserInput.xml

```
<MaxMailBoxSize>100</MaxMailBoxSize>
</UserGroup>
<Backup>
  <MaxParallelBackups>1</MaxParallelBackups>
  <BackupCapacity>1222222</BackupCapacity>
</Backup>
</UserInputs>
```

Exemplary ExchangeLayout.xml

```
<ExchangeLayout>
  <!-- Exchange Layout file generated by EEL -->
  . . . HEADER . . .
  <UserPriority>MaxCapacity</UserPriority>
  <OffHostBackup>>true</OffHostBackup>
  <ClusteringEnabled>>false</ClusteringEnabled>
  <Key>EX</Key>
  <RootMountPoint>ExchangeMountPoints</RootMountPoint>
  <EXServer Key="EX">
    <EXStorageGroup Key="EX">
      <Name>StorageGroup01</Name>
      <EXTransLogs Key="EX">
        <Name>StorageGroup01TransLogs</Name>
        <Size>1363148800</Size>
      </EXTransLogs>
    </EXStorageGroup>
    <EXDatabase Key="EX">
      <Name>TechSupport</Name>
      <EXEDB>
        <Name>TechSupportEDB</Name>
        <Size>3145728000</Size>
      </EXEDB>
      <EXSTM>
        <Name>TechSupportSTM</Name>
        <Size>3145728000</Size>
```

Exemplary Priority.xml

Installed by VERITAS Edition;  
Gives relative preferences for each type of storage to each type of Exchange volume.

```
<PriorityFile xmlns:xsi="http://www.w3.org/2000/10/XMLSchema-
instance"
xsi:noNamespaceSch
. . . HEADER . . .
  <Key>EX</Key>
  <PrioritySuggestions>
    <EXServer Key="EX">
      <EXStorageGroup>
        <DiskGroup>
          <EXTransLogs Key="EX">
            <DedicatedSpindleSetRequired>
              <MaxCapacity>>true</MaxCapacity>
              <Default>>true</Default>
              <HighPerfAvail>>true</HighPerfAvail>
            </DedicatedSpindleSetRequired>
            <PriorityConfig>
              <TypeOfRAID>
                <SWType>NORAI</SWType>
                <HWType>RAID0+1</HWType>
              </TypeOfRAID>
              <Priority>
                <Default>20</Default>
                <NumDisks ActualNumberOfSpindles="4">1
                </NumDisks>
                <TypeOfFS>NTFS</TypeOfFS>
              </PriorityConfig>
```



---

Exemplary SuggestedLayout.xml (storage configuration produced by the method)

---

```

<SuggestedLayoutFile xmlns:xsi="http://www.w3.org/2000/10/
XML.Schema-instance"
xsi:noNamespaceSchemaLocation=".\\StoreConfig.xsd">
  <!--SuggestedLayout.xml file generated by StorageConfigurator-->
  . . . HEADER . . .
  <RootMountPoint>ExchangeMountPoints</RootMountPoint>
  <VolumeLayout>
    <LocallyAttachedStorage>
      <DiskGroup ClusterDiskGroup="false">
        <Name>StorageGroup01DiskGroup</Name>
        <Volume>
          <Name>StorageGroup01TransLogs</Name>
        <MountPoint>F:\\ExchangeMountPoints\\StorageGroup01TransLogs
        </MountPoint>
          <FSType>NTFS</FSType>
          <Size>1363148800</Size>
          <RAIDType>
            <SWType>RAID1+0</SWType>
            <HWType>NORAID</HWType>
          </RAIDType>
          <ConsistsOf>
            <LU>
              <DisplayName>Harddisk7</DisplayName>
              <InternalID><![CDATA[SCSI\\DISK&VEN_
                IBM&PROD_DDRS-
                34560D&REV_
                </LU>
              <LU>
                <DisplayName>Harddisk8</DisplayName>
                <InternalID><![CDATA[SCSI\\DISK&VEN_
                IBM&PROD_DDRS-34560D&REV_
                </LU>
                . . . TOTAL OF 4 OF THESE FOR RAID1+0 (Mirrored
                Stripe)
              </ConsistsOf>
            <FilesOnVolume>
              <File>
                <FileName>StorageGroup01TransLogs</FileName>
                <ObjectType>EXTransLogs</ObjectType>
                <ParentObject>
                  <Name>StorageGroup01</Name>
                  <Type>EXStorageGroup</Type>
                </ParentObject>
              <PathName>F:\\ExchangeMountPoints\\StorageGroup01TransLogs\\
                StorageGroup01
              </File>
            </FilesOnVolume>
          </Volume>

```

---

### Clustering and Replication Embodiments

One embodiment of the invention provides configuration support for clustering (e.g., like Microsoft Cluster Service and VERITAS Cluster Service) and replication (VERITAS Volume Replicator) as well. In this embodiment, the “storage configurator” allows the user to select which servers on the network the user desires to replicate. Then the “storage configurator” gathers available disk information from all of the servers. The following example assumes replication of data from 3 sites (or 3 servers) around the world, these being NY, Seattle, and Tokyo). The user may then select the servers to configure for replication in one page of the wizard, then select the desired storage to use for each server on the next three pages, and then proceed through the rest of the wizard. At the end of the wizard, the storage configurator would produce three layouts, one for each server. The details of the volume layouts on each server may be different on each machine—depending on the storage available, but the volume names, sizes, disk group’s names, and application’s data objects names and volumes will all be the same. The user could then choose to implement the layouts on each

server. As a result, this embodiment operates to automatically configure clustering and/or replication.

Various embodiments may further include receiving, sending or storing instructions and/or data that implement a storage configurator in accordance with the foregoing description upon a carrier medium. Generally speaking, a carrier medium may include a memory medium as defined above, as well as transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as network and/or a wireless link.

Various modifications and changes may be made to the invention as would be obvious to a person skilled in the art having the benefit of this disclosure. It is intended that the following claims be interpreted to embrace all such modifications and changes and, accordingly, the specifications and drawings are to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. A system for determining a storage configuration for an application, the system comprising:

a computer system including a processor and memory;

and

one or more storage devices;

wherein the processor is operable to execute program instructions stored in the memory to:

display a graphical user interface for guiding a user through creation of the storage configuration;

receive user input to the graphical user interface, wherein the user input specifies one or more of: properties of the storage configuration or properties of the application, wherein the user input received to the graphical user interface specifies a mapping of one or more data objects associated with the application to one or more storage resources of the one or more storage devices; determine whether the specified mapping of data objects to storage resources is valid;

determine the storage configuration based on the user input received to the graphical user interface, wherein the storage configuration is determined using the specified mapping of data objects to storage resources if the mapping is determined to be valid;

wherein the storage configuration specifies a storage layout on the one or more storage devices for a plurality of data objects associated with the application.

2. The system of claim 1, wherein the processor is further operable to execute program instructions stored in the memory to:

implement the storage configuration, wherein said implementing the storage configuration comprises storing the plurality of data objects associated with the application on the one or more storage devices according to the storage layout specified by the storage configuration.

3. The system of claim 2, wherein the processor is further operable to execute program instructions stored in the memory to:

receive user input to the graphical user interface requesting implementation of the storage configuration;

wherein said implementing the storage configuration is performed in response to the user input requesting implementation of the storage configuration.

4. The system of claim 1, wherein the processor is further operable to execute program instructions stored in the memory to:

store the storage configuration in one of a file or a database.



29

5. The system of claim 4;  
 wherein said storing the storage configuration is performed at a first time;  
 wherein the processor is further operable to execute program instructions stored in the memory to:  
 retrieve the storage configuration at a second time subsequent to the first time;  
 re-display the graphical user interface; and  
 receive user input to the graphical user interface to modify properties of the storage configuration.

6. The system of claim 1,  
 wherein the processor is further operable to execute program instructions stored in the memory to:  
 display information indicating that the specified mapping of data objects to storage resources is not valid if the mapping is not valid.

7. The system of claim 1,  
 wherein said determining whether the specified mapping of data objects to storage resources is valid comprises determining whether storage resources specified by the mapping have sufficient storage capacity to hold the specified data objects.

8. The system of claim 1,  
 wherein said determining whether the specified mapping of data objects to storage resources is valid comprises validating the mapping against best practices of the application.

9. The system of claim 1,  
 wherein the one or more storage resources of the one or more storage devices comprise one or more volumes;  
 wherein the user input received to the graphical user interface specifies a mapping of one or more data objects associated with the application to one or more volumes.

10. The system of claim 1,  
 wherein the graphical user interface comprises a series of panels to guide the user through creation of the storage configuration in a wizard-like manner.

11. The system of claim 1,  
 wherein the plurality of data objects associated with the application include one or more of:  
 one or more files; and/or  
 one or more databases.

12. A method for determining a storage configuration for an application, the method comprising:  
 displaying a graphical user interface for guiding a user through creation of the storage configuration;  
 receiving user input to the graphical user interface, wherein the user input specifies one or more of: properties of the storage configuration or properties of the application, wherein the user input received to the graphical user interface specifies a mapping of one or more data objects associated with the application to one or more storage resources of the one or more storage devices;  
 determining whether the specified mapping of data objects to storage resources is valid;  
 automatically determining the storage configuration based on the user input received to the graphical user interface, wherein said automatically determining the storage configuration uses the specified mapping of data objects to storage resources if the mapping is valid;  
 wherein the storage configuration specifies a storage layout for a plurality of data objects associated with the application.

30

13. The method of claim 12, further comprising:  
 automatically implementing the storage configuration, wherein said implementing the storage configuration comprises storing the plurality of data objects associated with the application according to the storage layout specified by the storage configuration.

14. The method of claim 13, further comprising:  
 receiving user input to the graphical user interface requesting implementation of the storage configuration;  
 wherein said automatically implementing the storage configuration is performed in response to the user input requesting implementation of the storage configuration.

15. The method of claim 12, further comprising:  
 storing the storage configuration in one of a file or a database.

16. The method of claim 12,  
 wherein the method further comprises:  
 displaying information indicating that the specified mapping of data objects to storage resources is not valid if the mapping is not valid.

17. The method of claim 12,  
 wherein said determining whether the specified mapping of data objects to storage resources is valid comprises determining whether storage resources specified by the mapping have sufficient storage capacity to hold the specified data objects.

18. The method of claim 12,  
 wherein said determining whether the specified mapping of data objects to storage resources is valid comprises validating the mapping against best practices of the application.

19. The method of claim 12,  
 wherein the one or more storage resources comprise one or more volumes;  
 wherein the user input received to the graphical user interface specifies a mapping of one or more data objects associated with the application to one or more volumes.

20. The method of claim 12,  
 wherein the graphical user interface comprises a series of panels to guide the user through creation of the storage configuration in a wizard-like manner.

21. The method of claim 12,  
 wherein the plurality of data objects associated with the application include one or more of:  
 one or more files; and/or  
 one or more databases.

22. A computer-accessible memory medium for determining a storage configuration for an application, wherein the memory medium comprises program instructions executable by a processor to:  
 display a graphical user interface for guiding a user through creation of the storage configuration;  
 receive user input to the graphical user interface, wherein the user input specifies one or more of: properties of the storage configuration or properties of the application, wherein the user input received to the graphical user interface specifies a mapping of one or more data objects associated with the application to one or more storage resources of the one or more storage devices;  
 determine whether the specified mapping of data objects to storage resources is valid;  
 determine the storage configuration based on the user input received to the graphical user interface using the specified mapping of data objects to storage resources if the mapping is valid;



31

wherein the storage configuration specifies a storage layout for a plurality of data objects associated with the application.

23. The computer-accessible memory medium of claim 22, further comprising program instructions executable to: 5  
implement the storage configuration, wherein said implementing the storage configuration comprises storing the plurality of data objects associated with the application according to the storage layout specified by the storage configuration. 10

24. The computer-accessible memory medium of claim 22, further comprising program instructions executable to: store the storage configuration in one of a file or a database.

25. The computer-accessible memory medium of claim 22, 15  
wherein the computer-accessible memory medium further comprises program instructions executable to:  
display information indicating that the specified mapping of data objects to storage resources is not valid if the mapping is not valid. 20

26. The computer-accessible memory medium of claim 22, 25  
wherein the one or more storage resources comprise one or more volumes;  
wherein the user input received to the graphical user interface specifies a mapping of one or more data objects associated with the application to one or more volumes.

32

27. The computer-accessible memory medium of claim 22, wherein the graphical user interface comprises a series of panels to guide the user through creation of the storage configuration in a wizard-like manner.

28. The computer-accessible memory medium of claim 22, wherein the plurality of data objects associated with the application include one or more of:  
one or more files; and/or  
one or more databases.

29. The computer-accessible memory medium of claim 22, wherein said determining whether the specified mapping of data objects to storage resources is valid comprises determining whether storage resources specified by the mapping have sufficient storage capacity to hold the specified data objects.

30. The computer-accessible memory medium of claim 22, wherein said determining whether the specified mapping of data objects to storage resources is valid comprises validating the mapping against best practices of the application.

\* \* \* \* \*